

Oreo: Protecting ASLR Against Microarchitectural Attacks

Shixin Song

Massachusetts Institute of Technology
shixins@mit.edu

Joseph Zhang

Massachusetts Institute of Technology
jzha@mit.edu

Mengjia Yan

Massachusetts Institute of Technology
mengjiay@mit.edu

Abstract—Address Space Layout Randomization (ASLR) is one of the most prominently deployed mitigations against memory corruption attacks. ASLR randomly shuffles program virtual addresses to prevent attackers from knowing the location of program contents in memory. Microarchitectural side channels have been shown to defeat ASLR through various hardware mechanisms. We systematically analyze existing microarchitectural attacks and identify multiple leakage paths. Given the vast attack surface exposed by ASLR, it is challenging to effectively prevent leaking the ASLR secret against microarchitectural attacks.

Motivated by this, we present *Oreo*, a software-hardware co-design mitigation that strengthens ASLR against these attacks. *Oreo* uses a new memory mapping interface to remove secret randomized bits in virtual addresses before translating them to their corresponding physical addresses. This extra step hides randomized virtual addresses from microarchitecture structures, preventing side channels from leaking ASLR secrets. *Oreo* is transparent to user programs and incurs low overhead. We prototyped and evaluated our design on Linux using the hardware simulator gem5.

I. INTRODUCTION

Memory corruption vulnerabilities are some of the oldest security problems that continue to pose a serious security threat to modern systems [2], [72]. Among all the memory safety mechanisms proposed in the last few decades, Address Space Layout Randomization (ASLR) [36], [60], has shown to be effective in raising the barrier of attacks and has become one of the most prominently deployed mitigations in modern systems. ASLR works by randomly arranging the positions of code or data regions for the kernel or user-space applications. If the attackers cannot reliably determine the location of specific code or data, they will have difficulty carrying out control-flow hijacking attacks, such as return-oriented programming [62] and jump-oriented programming [10]. With ASLR, the attacker must perform an extra information disclosure step utilizing other existing vulnerabilities to leak ASLR secret before conducting their exploits.

However, ASLR has been defeated with various microarchitectural side channels. They pose a real threat, as we

witness an increasing number of such attacks being utilized in real-world software exploitations. For example, in 2017, a macOS kernel 0-day exploit [69] used the prefetch attack [31] to bypass ASLR. In 2022, a Linux kernel exploit (CVE-2022-42703) [38] also used side channels to bypass ASLR. In the same blog post, the authors stated that “KASLR is comprehensively compromised on x86 against local attackers, and has been for the past several years, and will be for the indefinite future.”

Among these microarchitectural-attack-assisted ASLR bypasses, a wide range of channels can be utilized, including TLBs, caches, and BTBs, using speculative execution, or even power-induced timing information [12], [13], [21], [22], [27]–[29], [31], [33], [37], [45], [47], [49], [52], [66], [76], [84]. Even worse, given the large attack surface, it seems that almost every newly discovered side-channel attack variant will likely become a new ASLR-bypassing attack vector. Given this phenomenon, two questions present themselves: (1) why has ASLR become such a fragile target for microarchitectural attacks, and (2) how can we secure ASLR to broadly block existing and potential future attack vectors?

Challenges. With a detailed investigation of existing ASLR bypasses using microarchitectural side channels, we find that the microarchitecture features that can be leveraged to use as a leakage channel are diverse and continue growing. Therefore, when protecting ASLR against side-channel attacks, addressing each individual channel or feature is not an appealing approach. For example, FLARE [13] narrowly focuses on closing a single channel, i.e., the address translation latency, while ignoring the abundant other side channels in modern processors, as well as other existing attack vectors. Instead, in this paper, we focus on blocking the leakage at the source by restricting the usage of the ASLR secret in both software and hardware.

Consider how the ASLR secret is used. ASLR shifts the location of a memory region by a *secret offset*. As such, the secret offset determines the virtual memory layout, i.e., which memory region is mapped and which is unmapped. The secret offset is also embedded in code and data pointers that are extensively used while executing a victim program. We systematically analyze and categorize microarchitectural ASLR bypasses into three leakage paths.

In the first leakage path, the attacker probes the virtual memory layout. This class of attacks relies on the fact that

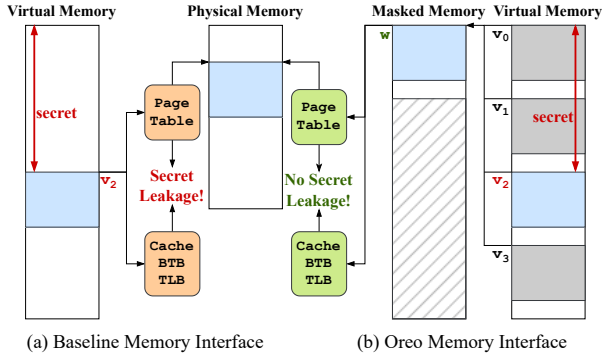


Figure 1: Overview of *Oreo*'s new memory interface

probing a mapped address versus an unmapped address resolves different microarchitectural side effects and thus distinct latency [12], [13], [27], [33], [37], [49], [66], [76], [78]. Additionally, these probe operations can be stealthy and do not cause system crashes, because they are either conducted under speculation or assisted with cache manipulation instructions such as prefetch.

Second, as the ASLR secret offset is embedded in code and data pointers, the attacker can leak the secret by monitoring the victim using its secret-dependent pointers to fetch instructions or perform loads and stores [21], [28], [29], [31], [45], [50], [52], [84]. These operations can result in distinguishable side effects on BTB, TLB, page table walker, caches, and DRAM.

Finally, the attacker can leak the victim's pointers by using the Spectre attack gadget and its variants. For example, the attacker may load the secret pointer into a register and then use it as the address of a load or a store instruction. There exists a substantial amount of work to block this leakage path, such as STT [82], NDA [77], InvisiSpec [79], and others [5], [6], [8], [15], [20], [23], [39], [40], [44], [48], [53], [58], [63]–[65], [67], [80], [81]. Unfortunately, existing mitigations targeting Spectre and its variants can only block the third leakage path, leaving the other two leakage paths unblocked. Moreover, many mitigations work by selectively delaying secret-dependent speculative execution. Such schemes are applicable to the backend of a processor (at the load/store queue) with moderate performance loss, but they are unappealing to the frontend of the processor, where delaying fetching instructions means leaving the rest of the processor seriously underutilized.

Oreo. The challenges discussed above motivate us to design *Oreo*, a software-hardware co-design scheme that secures ASLR against a wide range of microarchitectural side-channel attacks. Specifically, given a randomized virtual address, *Oreo* aims to protect selected randomized bits from being leaked via microarchitectural attacks following the first two leakage paths. We refer to these protected bits as *microarchitecture oblivious bits*.

The core of *Oreo* is a new memory interface, as shown in Figure 1. The left side of the figure shows the existing memory interface utilizing ASLR, mapping randomized virtual

addresses to physical addresses. A randomized virtual address, which embeds the secret ASLR offset, is used as input to look up entries in various hardware structures (caches, BTB, TLB, etc.) and critical software structures (the page table).

In contrast, *Oreo* introduces a new layer of memory, called *the masked memory*, sitting between the virtual and physical memory, shown on the right side of the figure. A masked address is constructed from a randomized virtual address with the microarchitecture oblivious bits being redacted. In this way, *Oreo* can map multiple virtual addresses to the same masked address and then map masked addresses to physical addresses using a modified page table. For example, in Figure 1, we show a valid region in the virtual address space starting with address v_2 . *Oreo* maps v_2 and another three invalid addresses v_0, v_1, v_3 to the same masked address w . All the hardware structures that used to use virtual addresses as input now switch to using masked addresses.

Oreo additionally changes the memory security check procedure. Given mapped and unmapped virtual addresses, which only differ in the microarchitecture oblivious bits, accessing them on *Oreo* will have the exact same microarchitectural side effects during speculation and only result in different architectural behaviors upon instruction commit time.

We further identify research challenges in selecting which bits to be protected by *Oreo*. In fact, not all the ASLR randomized bits can become microarchitecture oblivious bits. Moreover, it poses a security dilemma to obtain the entropy towards locating gadgets for both control-flow hijacking attacks and speculative execution attacks. We provide bits selection strategies to achieve entropy towards both attacks, and meanwhile to be adoptable by existing systems.

We prototype our design with support for protecting both kernel and user-space ASLR. We integrate the kernel changes on Linux 6.6, and implement our microarchitecture changes on the gem5 [9], [54] simulator. We show that our design introduces negligible performance overhead running the SPEC2017 IntRate benchmark [11] and the LEBench benchmark [85]. We provide a formal proof in a technical report [71] to show that *Oreo* achieves a non-interference property to prevent attackers from distinguishing virtual memory layouts with different secret offsets.

In summary, we make the following contributions:

- We systematically analyze existing microarchitectural attacks that leak the ASLR secret offset and classify them into three leakage paths.
- We propose a software-hardware co-design mitigation to prevent leaking selected bits of the ASLR secret, innovating a new memory interface.
- We prototype the software and hardware changes to support *Oreo* on Linux and the gem5 simulator.
- We provide security evaluation to show that *Oreo* prevents leakage of the ASLR secret, and our performance evaluation shows *Oreo* introduces negligible overhead.

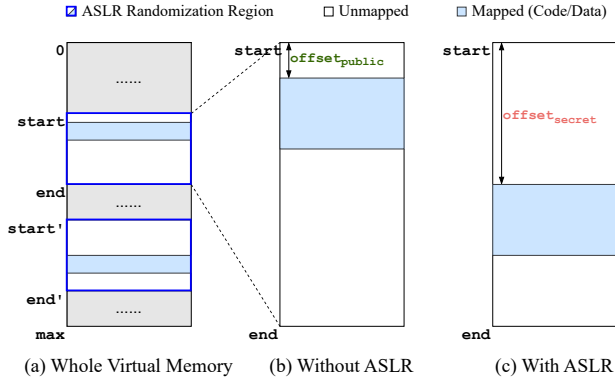


Figure 2: Coarse-grained ASLR. (a) shows the whole virtual address space with multiple randomization regions. (b) and (c) show memory content is loaded using a public offset when ASLR is disabled and a private offset when ASLR is enabled.

II. BACKGROUND

A. Address Space Layout Randomization

Address Space Layout Randomization (ASLR) [36], [60] is a widely deployed mitigation against memory corruption attacks. The idea is to randomly relocate code or data regions so that the attacker has difficulty determining the addresses for specific instruction gadgets to construct reliable code-reuse attacks. Figure 2 describes how coarse-grained ASLR works.

Given the full virtual address space, ASLR selects a memory region to perform its random re-location. In the whole virtual address space, multiple non-overlapping randomization regions exist for relocating different memory contents. For example, Linux kernel uses a region to hold kernel text and a different region for kernel modules. Figure 2(a) shows two randomization regions denoted as $[start, end)$ and $[start', end')$.

Within each randomization region, only a subset of address slots are used to hold code/data, thus considered valid addresses (the blue regions), while the rest are considered invalid. The length of the valid region is usually significantly smaller than the total size of the randomization region, which is necessary for achieving a reasonable amount of entropy. Address relocation shifts the valid region by an offset, which is the distance between the starting address of the valid region and the randomization region. This offset can be a public value when no ASLR is used (Figure 2(b)), and must be chosen secretly when ASLR is in place (Figure 2(c)).

B. Virtual Memory Systems

Modern systems support virtual memory for the purpose of process isolation, programmability, and hardware abstraction. With the virtual memory interface, the software does not directly operate on physical addresses backed by DRAM. Instead, the operating system abstracts DRAM by providing software with a large, contiguous, and unified virtual address

space, and introduces a layer of indirection to translate every virtual address to its mapped physical address.

Linux and many existing operating systems use page-based address translation. Taking the virtual page number (VPN) from a virtual address, a *page table walk* translates the VPN to its corresponding physical page number (PPN). Modern systems use hierarchical page tables to store the page table in a space-efficient manner. The CPU looks up virtual addresses by traversing a tree structure from root to leaf, requiring multiple memory accesses. Microarchitecture structures, such as translation lookaside buffers (TLBs) and page table caches, can buffer recently accessed translations to accelerate this procedure.

The virtual memory system enforces security checks upon every memory access leveraging protection bits embedded in page table entries (PTEs). The check involves checking whether a virtual address is mapped or not. Accessing an unmapped address leads to looking up an invalid PTE entry and thus results in a page fault. The virtual memory security check additionally examines whether the access to the page has the correct permissions from an appropriate privilege level.

C. Microarchitectural Side Channel Attacks

A microarchitectural side channel involves information leakage from a victim’s security domain to an attacker’s security domain. The attacker exploits visible side effects of the execution of instructions whose behaviors are secret-dependent. We call such instructions *transmitters*. The transmitter leaves side effects by modifying the states or occupancy of various microarchitecture structures, such as caches [41]–[43], [51], [56], [68], TLBs [28], [45], BTB [16], [21], and Network-on-Chips [19]. Furthermore, recent microarchitectural attacks exploit power-induced timing leakage [75].

Speculative execution attacks, also referred to as transient execution attacks, are a class of information leakage attacks where attackers exploit the side effects of *transient* instructions. A transient instruction is an instruction that is speculatively executed on an out-of-order core but is later squashed due to misspeculation. High-profile speculative execution attacks include Meltdown [51], Spectre [42], and its variants [41], [43], [56], [68].

Given that modern processors have been aggressively optimized, an increasing number of microarchitectural side channels and speculative execution features have been revealed in the last few years. We envision this trend will continue. Furthermore, as the arms race continues, due to performance overhead and hardware costs, it is unlikely that the industry will commit to delivering processors with comprehensive mitigations in the near future. Hence, non-complete solutions that only block certain leakage channels are insufficient. This paper presents a mitigation that blocks microarchitectural-attack-assisted ASLR bypasses. Our scheme works even if the processor exhibits vulnerable microarchitectural leakage channels and speculation features.

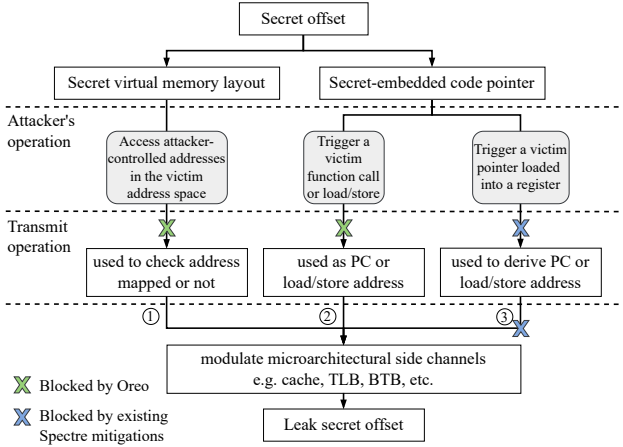


Figure 3: Three leakage paths that leak ASLR secret offset

III. UNDERSTANDING ASLR BYPASSES USING MICROARCHITECTURAL ATTACKS

ASLR can be bypassed by leaking the secret offset using either software attacks or microarchitectural attacks. In this paper, we focus on microarchitectural-attack-assisted ASLR bypasses because most of these attacks do not need to exploit software bugs and can universally work with a wide range of commercial processors. In fact, these attacks have been put into use in real-world software exploits [38], [61], [69].

We systematically analyze the past microarchitectural-attack-assisted ASLR bypasses and find there exist multiple possible leakage paths. As the ASLR secret is spread throughout the system, it unavoidably exposes a large attack surface. Specifically, the secret offset determines the virtual memory layout, meaning it determines which regions are mapped and which are not. The secret offset is also embedded in victim pointers, which are extensively used by the processor to fetch/load/store when executing a victim program.

Figure 3 summarizes these attack surfaces. From top to bottom, we show where the secret offset is located (the top row), what attacker operations (the second row) can be used to trigger a secret-dependent transmission operation (the third row), and what microarchitectural side channels the transmission operation can modulate (the bottom row). Overall, the figure shows three leakage paths. Moreover, the microarchitectural channels (structures) that can be used to leak the secret are diverse, indicating blocking each single side channel is not a promising strategy to defend against ASLR bypass attacks. A detailed enumeration of existing attacks is provided in Appendix B, including the leakage paths they take and the utilized side channels.

Virtual Memory Layout Probing. The first leakage path (the left path in Figure 3) shows an attacker probes the virtual memory layout to figure out which region is mapped and which is unmapped [12], [13], [27], [33], [37], [49], [66], [76], [78]. The attacker’s operation involves triggering the processor to access an attacker-controlled address in the

victim address space. The processor needs to consult several microarchitecture structures and perform a page table walk if needed to determine whether the address is mapped or not, resulting in distinguishable microarchitectural side effects. Usually, accessing unmapped addresses will result in system crashes or exceptions. However, microarchitectural attacks can be stealthy since a clever attacker can suppress the crashes using speculation and other tricks, such as cache manipulation instructions (e.g., prefetch) and Intel TSX [35].

Multiple existing attacks fall into this category. For example, as shown in the following code snippet, DrK [37] attacks kernel ASLR by probing each page in the kernel’s randomization region from the user space and checking whether the page is mapped or not.

```

1 // guess_addr is a kernel address, controlled by the attacker
2 for (guess_addr = start; guess_addr < end;
3     guess_addr += page_size) {
4     // Step 1: Access guess_addr, suppress crashes with Intel
5     // TSX. TLB caches transformation if guess_addr is valid.
6     tsx_probe(guess_addr);
7     // Step 2: Access guess_addr again, measure page fault
8     // latency. Latency is low if the translation is in TLB.
9     page_fault_latency = tsx_probe(guess_addr);
10    if (page_fault_latency < threshold) { break; }
11 }

```

For each page, the attacker probes (e.g., issue a load) to a virtual address (denoted as `guess_addr`) in that page two times. In both instances, the attacker uses Intel TSX to suppress page faults caused by failed permission checks. Specifically, instead of informing the OS and causing a real crash, Intel TSX invokes a user-specified exception handler when a page fault occurs [35]. At the first probe, a page table walk (PTW) is triggered. The PTW inserts a PTE entry into the TLB if the tested page is mapped and leaves the TLB unchanged otherwise. The attacker then issues the second probe to test whether the address translation for `guess_addr` is cached in the TLB. This test is performed by measuring how long it takes for the second probe to trigger a page fault that is then caught by the attacker’s exception handler.

As shown above, attacks in this leakage path rely on distinguishing between mapped and unmapped addresses based on their microarchitectural side effects. In addition to using TLB behaviors, prior work [12], [13], [66], [76] also discovered that pipeline behaviors differ for mapped and unmapped addresses. Taking the Data Bounce attack [12], [66] as an example, they exploit the timing difference introduced by the store-to-load forwarding scheme, which is only triggered when the store address is mapped.

Leaking Pointers as Addresses. The middle path in Figure 3 shows how the secret offset embedded in a victim pointer can be leaked when it is used as an *address* during the victim’s execution. Specifically, the pointer is used as the program counter or a load/store address. Unlike the first leakage path, this attack vector does not require speculation or any crash suppression schemes. For example, to use a victim address as the program counter, the attacker only needs to trigger the victim to do a function call *non-speculatively*. When accessing a secret-dependent address, various microarchitecture structures

will be modulated, including BTBs (Jump Over ASLR [21]), TLBs (TLBBleed [28], TagBleed [45], the Prefetch attack [31], and EntryBleed [52]), and page table walkers (AnC [29] and Binoculars [84]).

We provide an example below to illustrate the AnC attack [29], which uses cache Prime+Probe to leak the kernel ASLR secret from user space. After resetting the cache states, the attacker simply makes a system call, which triggers the victim (i.e., the kernel in this example) to fetch instructions using secret-dependent virtual addresses and the processor will trigger the page table walker (PTW) to translate these virtual addresses. Note that, as the PTW uses secret bits of the virtual addresses to index into page tables and modulate the caches, the attacker can leak the secret offset by monitoring cache states.

```

1 // Step 1: reset cache states
2 prime_the_whole_cache();
3 // Step 2: trigger the victim to fetch mapped virtual
4 // addresses and the PTW process to translate these
  → addresses.
5 syscall();
6 // Step 3: Use cache Prime+Probe to figure out which cache
  → set
7 // has been modulated by the PTW. "hit_set" reveals
8 // selected bits in the virtual addresses used by PTW
9 hit_set = probe_every_cache_set();

```

Leaking Pointers as Data. The final leakage path (shown as the right path in Figure 3) describes how the secret embedded victim pointers can be leaked as data. Different from the previous leakage path, the attacker needs to leverage a memory corruption vulnerability or transient out-of-bound memory access to load a victim pointer into a register and then use the secret bits in the pointer to compute an address. This derived address is then leaked via side channels.

This leakage path usually requires a classic Spectre attack gadget or its variants [41]–[43], [56], [68]. Such gadgets pose a serious threat as they can be exploited to leak arbitrary data in the victim’s address space, not just pointers. As such, extensive work has been proposed to mitigate this threat. For example, STT [82] and NDA [77] delay speculative execution of transmission instruction (e.g., load/store) if their operand holds speculative data. Many other prior works [5], [6], [8], [15], [20], [23], [39], [40], [44], [48], [53], [58], [63]–[65], [67], [79]–[81] intend to hide side effects of speculative execution on the cache hierarchy and other structures.

IV. THREAT MODEL

We follow the threat model of microarchitectural side channel attacks, where the attacker and the victim reside in different security domains. This setup includes the case when the attacker and the victim are two user-space processes, or the victim is a privileged software, such as an operating system kernel and hypervisor, while the attacker is a user-space application. This also applies when the victim is an enclave program while the attacker is privileged software. The attacker and the victim execute on the same machine, sharing various microarchitecture structures. We broadly consider timing-based side channels due to resource contention and

speculation. Our threat model does not specifically consider power-induced timing side channels [49], but we elaborate on how our scheme can help future mitigations in Section VIII.

We set out to block the first two ASLR leakage paths in Figure 3, indicated by the two green crosses. We propose *Oreo*, a software-hardware co-design scheme to prevent leakage through direct virtual memory layout probing and ensure secret-embedded pointers, when used as program counters or load/store addresses, remain indistinguishable to attackers. *Oreo* is highly practical and can be adopted in real systems.

As a side note, substantial prior work [5], [6], [8], [15], [20], [23], [39], [40], [44], [48], [53], [58], [63]–[65], [67], [77], [80]–[82] has attempted to address the data leakage path (the third path). These schemes can be used in complementary with *Oreo* to serve as a comprehensive defense solution for achieving security goals beyond mitigating ASLR bypasses.

V. DESIGN OF *Oreo*

A. Overview

The goal of *Oreo* is to protect selected randomized bits in virtual addresses from being leaked via microarchitectural attacks in the first two leakage paths in Figure 3. We refer to these bits as *microarchitecture oblivious bits* or *protected bits* for short. The protected bits in a valid randomized virtual address concatenated with trailing zeros form the secret offset. Without ambiguity, we refer to the secret offset protected by *Oreo* as $\text{offset}_{\text{oreo}}$. For example, given a valid randomized address $0xFFAB12340$, if configuring the protected bits as bits 20 to 27, then the secret $\text{offset}_{\text{oreo}}$ is $0xAB00000$.

How to Protect Microarchitecture Oblivious Bits? We introduce a new memory interface, with an extra layer of masked address space that sits between the virtual address space and the physical address space. A masked address is mapped from a virtual address with the microarchitecture oblivious bits redacted. As shown in Figure 4, multiple virtual addresses with different microarchitecture oblivious bits are mapped to the same masked address. *Oreo* uses masked addresses to build and traverse page tables and access various microarchitecture structures, such as BTBs and TLBs. This scheme ensures accessing a virtual address results in microarchitectural side effects independent from its protected bits.

In addition, *Oreo* changes the memory address security check flow. Using masked addresses makes unmapped and mapped virtual addresses that only differ in microarchitecture oblivious bits have exactly the same microarchitectural side effects, but we eventually need to distinguish between them to raise exceptions upon illegal memory accesses. *Oreo* performs this check at the commit time of instructions, forcing the distinguishability between these addresses to happen only at the architectural level rather than the microarchitectural level.

How to Choose Which Bits to Protect? There exist several constraints in choosing the bits to be protected by *Oreo*. For example, *Oreo*’s protection faces a security dilemma. On the positive side, *Oreo* prevents leaking $\text{offset}_{\text{oreo}}$ using microarchitectural attacks, strengthens the security of ASLR, and raises the barrier against control-flow hijacking attacks.

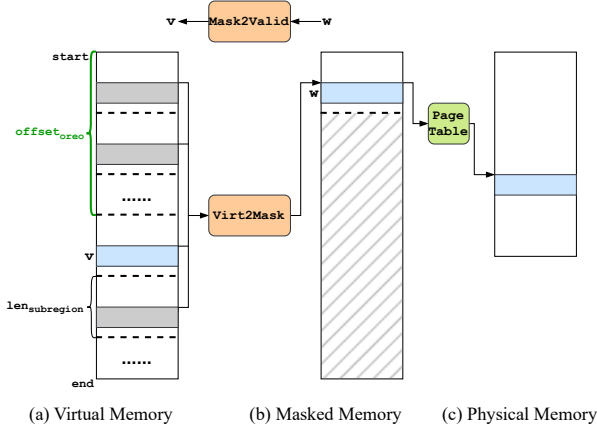


Figure 4: Mapping between a virtual address v and a masked address w using *Oreo*'s memory interface.

However, on the negative side, an attacker who aims to perform speculative execution attacks can execute a Spectre-like gadget under speculation without knowing the secret $\text{offset}_{\text{oreo}}$. This is because accessing addresses only differing in the *Oreo* protected bits speculatively results in the same microarchitectural side effects.

We present two strategies to determine which bits to be protected by *Oreo*. A naive strategy chooses to protect part of the default baseline ASLR randomized bits, and an enhanced strategy introduces additional entropy and protects bits that do not overlap with the randomized bits in the baseline.

B. Protecting Microarchitecture Oblivious Bits

Masked Address Space. *Oreo* introduces the masked address space and maps multiple virtual addresses to the same masked address. Figure 4 shows the notation we use to describe the masked address space layout. Given a memory range to be used for ASLR denoted as $[\text{start}, \text{end})$, we divide this range into multiple subregions with equal size denoted as $\text{len}_{\text{subregion}}$, marked by the dotted lines in Figure 4(a). Addresses in the i -th region have their microarchitecture oblivious bits equal to i , and the goal is to prevent these bits from being leaked.

Oreo maps each subregion to the very first subregion in the masked address space. Given a valid address v , the corresponding masked address can be calculated as:

$$\text{Virt2Mask}(v) = ((v - \text{start}) \bmod \text{len}_{\text{subregion}}) + \text{start}.$$

We note two things from the formula above. First, the calculation of the virtual-to-masked memory mapping, as well as the masked-to-physical mapping, are independent of microarchitecture oblivious bits. Second, it is simple enough to be implemented efficiently in hardware with low cost.

After the address conversion is done, *Oreo* needs to set up page tables to map masked addresses instead of virtual addresses. We discuss required kernel changes in Section VI-A. In addition, any microarchitecture structures that use virtual

addresses as indices will be modified to use masked addresses, as detailed in Section VI-B.

Security Check Flow. *Oreo* forces the virtual addresses, as long as they map to the same masked address, to be only distinguished at the architectural level after the instruction commits. In other words, accessing unmapped addresses eventually results in exceptions upon instruction commits. We clarify how *Oreo*'s security check flow interacts with the existing virtual memory security check.

Oreo's security check flow involves two steps. The first step is performed during speculation and is exactly the same as the existing virtual memory security check, with the only difference being performing the check upon masked addresses rather than virtual addresses. Recall from Section II-B; the existing virtual memory security check examines whether the accessed page is mapped as well as whether the access has the correct read/write/execute permissions from an appropriate privilege level. This check is performed during TLB accesses or page table walks. *Oreo* applies this check on masked addresses during speculation so that this check is independent from the microarchitecture oblivious bits.

The second step of the security check happens at the commit time. The task is to check whether the virtual address has the correct microarchitecture oblivious bits. This can be done by reconstructing the valid virtual address and checking whether the virtual address used by the committing instruction is equal to the valid virtual address. The following formula reconstructs the valid virtual address from a masked address w .

$$\text{Mask2Valid}(w) = \text{offset}_{\text{oreo}} + w.$$

This formula indicates that we need to store the secret $\text{offset}_{\text{oreo}}$ somewhere to be readily used upon instruction commit time. In our implementation, we store this information in the unused bits in page table entries [35], which we detail in Section VI-A.

Blocking Leakage Paths. We show below how *Oreo* blocks the first two leakage paths using concrete examples. Formal reasoning is presented in a technical report [71]. Consider the following example where a program is loaded at a random address $0\text{xFFAB}12340$. If *Oreo* intends to protect $\text{offset}_{\text{oreo}} = 0\text{xAB}00000$, the corresponding masked address can be computed as $0\text{xFF}0012340$.

In the first leakage path, the attacker probes the victim's virtual address space by issuing speculative memory access operations targeting different victim addresses and expecting to distinguish between mapped and unmapped addresses based on their microarchitectural side effects. For example, when performing the Double Page Fault attack [33] and Code Region Probing attack [27] on the baseline insecure system, if the attacker accesses any address with incorrect microarchitecture oblivious bits (e.g., $0\text{xFFAA}12340$), the TLB will not be filled for this invalid address translation, and the caches will not be filled with such an invalid address. This differs from the microarchitectural behaviors when accessing the valid address.

Using masked addresses allows *Oreo* to have the exact same microarchitectural effects when accessing addresses that differ

in microarchitecture oblivious bits for two reasons. First, on *Oreo*, virtual addresses that fall within the randomization will be converted to masked addresses before address translation and memory accesses. Both the invalid address `0xFFAA12340` and the valid one `0xFFAB12340` will be converted to the same masked address `0xFF0012340`. In both cases, the TLB will be filled with a valid address translation, and a valid cache entry located using the translated physical address will be inserted into the cache hierarchy. Second, *Oreo* delays the check until the commit time to determine whether an address is mapped or not. Recall that, as the attacker aims to suppress potential exceptions using pipeline squashes, none of the attacker’s probing operations will reach the commit stage. As a consequence, they follow the same security check under speculation and their microarchitectural effects will not be distinguishable on *Oreo*.

Reasoning about how *Oreo* blocks the second leakage path is straightforward. Recall that in the second leakage path, the attacker triggers a victim functional call to make the victim branch to a valid virtual address (e.g., `0xFFAB12340`). In the insecure baseline, the secret bits `0xAB` are used as part of the address to locate entries in TLBs, BTBs, and caches. While, *Oreo* redacts the secret bits from the virtual address to obtain the secret-free masked address `0xFF0012340`. Furthermore, all the microarchitecture structures that used to take virtual addresses as input now switch to using masked addresses, leading to no side effects related to the secret `offsetOreo`, so that *Oreo* successfully blocks the second leakage path.

C. Choosing Bits to Protect

Constraints. Before presenting our strategies in choosing the bits to be protected by *Oreo*, let’s first understand what constraints we have.

First, the least significant microarchitecture oblivious bit is constrained by the subregion size used by *Oreo*. Recall from Figure 4, *Oreo* divides the memory range to be used by ASLR into multiple equally-sized subregions and maps these subregions to a single subregion in the masked address space. It is required that only one of the subregions in the virtual address space is valid. As such, we need to ensure the subregion size $\text{len}_{\text{subregion}}$ to be large enough to hold the code or data that is being relocated. For example, the Linux kernel text is $2^{25} = 32$ MB, and thus the least significant bit we can pick to be protected by *Oreo* is bit 25. The baseline ASLR does not have this constraint since the valid region can be shifted at the page granularity. The least significant bit to be randomized is bit 12 if using 4 KB pages or bit 21 if using 2 MB pages.

Second, we face a security dilemma. *Oreo* makes mapped and unmapped addresses to be indistinguishable at the microarchitectural level if they only differ in microarchitecture oblivious bits. The problem is that if attackers aim to conduct speculative execution attacks, they can utilize Spectre-like gadgets without knowing the correct microarchitecture oblivious bits in the gadgets’ randomized virtual addresses. The attack can succeed because accessing virtual addresses

with incorrect microarchitecture oblivious bits has the same microarchitectural effects during speculation as accessing the one with the correct bits.

If we increase the bits to be protected by *Oreo*, we increase the system’s resilience against ASLR bypass attacks and thus make control-flow hijacking attacks, such as ROP attacks, more difficult. However, we may face the risk of decreasing the time it takes an attacker to locate and exploit Spectre gadgets. Therefore, we need to choose which bits to protect carefully to optimize for the entropy we can achieve against both attacks. **Bits Selection Strategies.** Figure 5 lists the ASLR randomized bits and microarchitecture oblivious bits in four different setups, including the default baseline ASLR used in Linux, a naive *Oreo* bits selection strategy, an enhanced baseline with additional randomized bits, and an enhanced *Oreo* selection strategy. In the baseline ASLR setup used by Linux (top row in the figure), the n bits in the virtual address, colored in blue, are randomized, and the lower k bits are not randomized.

The naive strategy is to choose part of the ASLR secret bits as the microarchitecture oblivious bits. For example, we can choose the higher m bits (colored in green) of the ASLR secret bits, so the least significant bit protected by *Oreo* is the $(k+n-m)$ th bit. Due to the subregion size constraint discussed before, we have to ensure m is small enough so that we have 2^{k+n-m} to be at least as large as the valid region size. The concern with this strategy is when the value of m is too small, *Oreo* has limited entropy towards mitigating ASLR bypasses. Besides, it is not effective in addressing the speculative execution security dilemma discussed above either.

We address the above constraints by proposing an enhanced bits selection strategy, which introduces extra entropy in addition to the entropy of the default ASLR. For a fair comparison, we present an enhanced baseline (the third row) and the enhanced *Oreo* strategy (the bottom row) in Figure 5. The enhanced baseline additionally randomizes the higher m bits in the virtual address to the left of the n bits already randomized by the default ASLR. *Oreo*’s enhanced bits selection strategy randomizes the same $m+n$ bits as the enhanced baseline, but protects the higher m bits (marked as green). In this case, we no longer need to be concerned with the subregion size constraint since the least significant bit protected by *Oreo* is not smaller than $(k+n)$. 2^{k+n} is the default ASLR randomization region and is deemed to be larger than the valid region size.

Entropy Analysis. We summarize the entropy comparison of these schemes in Figure 5. We show the entropy of each scheme against locating gadgets to be used in code reuse attacks and speculative execution attacks. We then compare the original entropy of each scheme with their remaining entropy after the attacker performs any of the microarchitectural-based ASLR bypasses in the first two leakage paths in Section III. Note that blocking speculative execution attacks is out of the scope of our threat model; rather, we want to ensure *Oreo* does not make this type of attack easier.

We begin by clarifying the security implications of the original entropy. It is desired to retain the original entropy

| | | | | Code Reuse Attacks | | Speculative Execution Attacks | |
|-------------------------|--|--|--|--------------------|--------------------|-------------------------------|--------------------|
| Locating Gadgets for... | | | | Original Entropy | Remaining Entropy* | Original Entropy | Remaining Entropy* |
| Baseline-Default | | | | n | 0 | n | 0 |
| <i>Oreo</i> -Naive | | | | n | m | $n - m$ | 0 |
| Baseline-Enhanced | | | | $n + m$ | 0 | $n + m$ | 0 |
| <i>Oreo</i> -Enhanced | | | | $n + m$ | m | n | 0 |

Figure 5: Protected bits selection strategies and their corresponding entropy. “Remaining entropy*” refers to “remaining entropy after ASLR bypasses using the first two leakage paths”

of *Oreo* to match the corresponding baseline ASLR. In this way, we force the attacker to pay extra effort to perform ASLR bypass attacks. However, further increasing the original entropy does not help strengthen the security of ASLR against microarchitectural attacks. Several attacks in leakage path ②, such as the AnC attack [29] and Binoculars [84], directly leak all ASLR randomized bits. Importantly, their leakage time is independent of the number of bits to be leaked. As shown in Figure 5, though the enhanced baseline has higher original entropy than the default baseline, they both have no resilience against microarchitectural-attack-assisted ASLR bypasses and have 0-bit remaining entropy against locating gadgets for both code reuse attacks and speculative execution attacks.

With the naive *Oreo* bit selection strategy, we obtain m -bit entropy against gadget detection for code reuse attacks. The original entropy of gadget detection for speculative execution attacks is reduced from n bits to $n - m$ bits since the attacker does not need to know the microarchitecture oblivious bits. Using the enhanced bits selection scheme, we have m -bit remaining entropy against gadget detection for code reuse attacks and n -bit original entropy against gadget detection for speculative execution attacks.

To summarize, introducing extra randomized bits that are not protected by *Oreo* does not gain security against microarchitectural attacks. Even though the enhanced baseline provides m more bits of original entropy compared to the default baseline and the *Oreo* enhanced strategy, it increases little security guarantee against speculative execution attacks. Overall, the *Oreo* enhanced bits selection strategy achieves the best security property, effectively increasing the barrier against control-flow hijacking attacks and retaining the barrier against speculative execution attacks.

Feasibility and Linux Prototyping. It is feasible to adopt the enhanced *Oreo* bit selection strategy in existing systems. We have implemented this strategy in our Linux prototype for kernel text, kernel modules, and user-space programs.

For kernel text, Linux’s default configuration relocates the code within a 1GB region using 2MB as the relocation alignment. As such, the default ASLR randomizes bits 21 to 29, providing a 9-bit entropy. Similarly, ASLR relocates kernel modules within a 1GB region that is consecutive to the kernel text randomization region and uses 4KB as the alignment size, so bits 12 to 29 are randomized. To reserve large enough memory for kernel modules, the default ASLR allows 1024 possible offsets for relocation, providing 10-bit entropy.

Following the enhanced strategy in Figure 5, we need to randomize additional bits higher than the ASLR secret bits. According to the Linux kernel memory management documentation [3], there exists a consecutive 444 GB unused region that can serve our randomization goal. We use this region for both kernel text and modules, so each of them can use 222 GB. We configure *Oreo* to protect bits 31 to 38 (8 bits in total) for both the kernel text and kernel modules.

We can also apply the enhanced selection strategy to kernel data regions and user-space memory. For example, the default user-space ASLR can use the whole user-space virtual address space as the randomization region with a granularity of 4 KB and it provides a high entropy of 28 bits. To maximize applicability, we do not want to reduce the available virtual memory size for the user space. Therefore, we choose non-canonical bits as microarchitecture oblivious bits. Considering a memory system using 4-level page tables, a canonical address is derived by taking a 48-bit virtual address and sign-extending it to form a 64-bit address. Bits 48 and above are not used in the baseline, so using them as microarchitecture oblivious bits will not affect available memory size.

D. Further Increasing Entropy of *Oreo*-Protected Bits

So far, we have shown *Oreo* works with coarse-grained ASLR, where the addresses from the same valid region (the blue region in Figure 4(a)) share the same secret offset. However, one limitation of coarse-grained ASLR is that leaking one pointer breaks the whole defense. We now examine how to further increase the entropy of *Oreo*-protected bits.

Working with Existing FGASLR. The existing fine-grained ASLR randomizes the memory layout by (1) shuffling the order of functions inside the program and (2) relocating the program by a random offset (same as the coarse-grained ASLR). *Oreo* can only protect the random offset in (2) but cannot prevent leaking the order of functions after shuffling.

Supporting Page-Granularity ASLR. An appealing feature of *Oreo* is its capability to conveniently support page-granularity ASLR to improve its safety level. Figure 6 shows a configuration where we have four pages in the valid region. Figure 6(a) describes when using coarse-grained ASLR, the four pages use the same $\text{offset}_{\text{oreo}}$ to be relocated to the third subregion in the virtual address space. Alternatively, Figure 6(c) shows the case when the four pages are relocated to different subregions in the virtual address space while still mapped to non-overlapped pages in the masked address space.

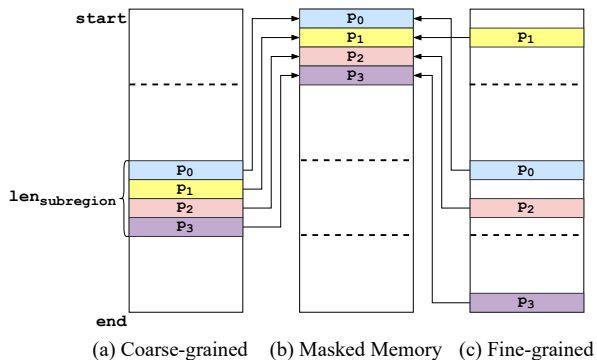


Figure 6: Compare *Oreo*'s randomized memory layout for coarse-grained ASLR (left) and page-granularity ASLR (right)

We now analyze the `Virt2Mask` and `Mask2Valid` functions to highlight that minimal changes are needed to make *Oreo* support page-granularity ASLR. First, the coarse-grained and the fine-grained schemes use the same `Virt2Mask` function. Second, the `Mask2Valid` function only differs slightly as the page-granularity scheme needs to configure different `offsetoreo` for different pages, which means storing different `offsetoreo` values in the corresponding PTE entries.

Configuring *Oreo* is handy if page-granularity ASLR is in place. However, we acknowledge that there exist engineering challenges to implement a proper code relocater for such a randomization scheme in software.

E. Limitations

While *Oreo* offers substantial improvements in ASLR security, it also has several limitations that must be considered. In this section, we summarize these limitations to help understand the trade-offs involved in applying *Oreo*.

Security Dilemma. In Section V-C, we have detailed the security dilemma problem and how we address it by carefully selecting bits to protect. Basically, on the one hand, *Oreo* successfully protects ASLR randomized bits against the first two leakage paths and makes control-flow hijack attacks more difficult. On the other hand, the protected bits cannot prevent attackers from utilizing Spectre gadgets since they can speculatively execute the gadgets without knowing these bits. When applying *Oreo*, we need to take this dilemma into consideration and carefully choose bits to protect so that the overall system achieves a satisfactory security guarantee against both attacks.

Using Non-Canonical Bits or Canonical Bits. When applying *Oreo* to user-space programs, we choose microarchitecture oblivious bits from non-canonical bits. We see no problem in applying such a scheme in systems using 4-level page tables where we implement our prototype, but it may introduce limitations due to reduced non-canonical bits in other system configurations.

For example, given a system using 5-level page tables and 57-bit virtual addresses, *Oreo* can use no more than 7 bits as

microarchitecture oblivious bits. The non-canonical bits might also be used for other purposes such as ARM PAC [34], which further limits the entropy. To achieve higher entropy, we can use canonical bits as microarchitecture oblivious bits for user-space programs, following a similar configuration used by *Oreo* for kernel text and modules.

However, using canonical bits faces another constraint due to the program's code or data sizes. As analyzed in Section V-C, the least significant microarchitecture oblivious bit is constrained by the subregion size (`lensubregion`) used by *Oreo*, which must be large enough to hold the relocated code or data. Different programs have varying code and data sizes. Enforcing a uniform configuration for all user-space programs to accommodate the largest programs would unnecessarily limit the entropy for smaller programs. Instead, *Oreo* can adopt a more flexible approach. Specifically, `lensubregion` can be adjusted based on the size of each program, allowing *Oreo* to protect more bits and achieve higher entropy for smaller programs.

Using PTE Bits. As *Oreo* stores the protected bits in PTEs, the maximum entropy of the protected bits is limited by the number of unused bits in PTEs. For example, PTE bits can be used by other system software and security mechanisms such as MPK [4]. In our prototype, we have already considered these factors and used PTE bits that do not overlap with MPK bits and software-available bits. Specifically, our prototype (Section VI-A) uses leaf PTEs to store microarchitecture oblivious bits, which allows at most 9 bits for the kernel space and 5 bits for the user space. If future software uses more leaf PTE bits, *Oreo* can use the unused bits in other levels of PTEs.

VI. IMPLEMENTATION DETAILS

A. Software Changes

Oreo's software changes require modifying the page table to use masked addresses. We prototype the software changes in Linux kernel. Our current prototype supports three types of memory regions: kernel text, kernel modules, and user-space programs. It is also feasible to adopt *Oreo* for other memory regions, including kernel data regions.

For all the randomization regions, we set up the page table to map masked addresses to physical addresses and record the secret `offsetoreo` in PTE entries. Our implementation is compatible with the existing Linux kernel implementation. The implementation for the three regions slightly differs on 1) which bits are selected to be microarchitecture oblivious bits, and 2) when a page mapping is set up.

Kernel Text. Linux kernel uses 21 to 29 as the ASLR randomized bits. Our prototype uses the enhanced bits selection strategy (Section V-C) to additionally randomize and protect bits 31 to 38, i.e., the microarchitecture oblivious bits. We store these bits at the PTE entry for each kernel text page.

Several Linux-specific implementation details are worth mentioning. First, Linux sets up the page table for kernel text at boot time. Therefore, our prototype directly integrates our changes at the boot code. Second, since our prototype uses 8 microarchitecture oblivious bits and the x64 architecture

reserves 9 unused bits in leaf PTEs [46], our changes to the PTE entries are compatible with existing implementations without affecting the current page tables’ functionality.

Kernel Modules. Similar to kernel text, we choose bits 31 to 38 as the microarchitecture oblivious bits for kernel modules.

We note one implementation detail specific to the current implementation of Linux. In our prototype, we use the same $\text{offset}_{\text{oreo}}$ for kernel modules and kernel text. This differs from the baseline ASLR, where the kernel modules can have different entropy from kernel text. The reason is that the current Linux implementation requires the kernel text and modules located in a 2 GB consecutive region, so bits 31 to 38 (the microarchitecture oblivious bits) in their virtual addresses need to be the same. If we want to support different entropy for kernel text and modules, we will need to relax this constraint by leveraging prior work, such as Adelle [59], which allows relocating kernel modules in the whole 64-bit address space.

We make the following changes to the page setup procedure for kernel modules. Linux allocates memory for a module when the module is loaded into the kernel. The virtual memory allocation triggers page table entry setup. Additionally, Linux builds a red-black tree to manage memory for these modules. In the default implementation, each module’s randomized virtual base address is used as the key to construct and search the tree, which introduces a side-channel vulnerability. We change the kernel to allocate and manage memory using masked addresses. In addition to properly setting the page table entries, we also use modules’ masked base addresses to build the red-black tree.

User-Space ASLR. As discussed in Section V-C, we use the non-canonical bits as microarchitecture oblivious bits for user-space applications. Theoretically, we have 16 non-canonical bits. In our prototype, we use bits 48 to 52, providing 5-bit entropy, which matches the number of unused bits in user-space leaf PTEs (rather than 9 unused bits in kernel leaf PTEs). We note that this is an engineering decision made for convenience. It is feasible to increase the entropy by storing the extra microarchitecture oblivious bits in the PTEs of the other levels of page tables. As each user process has its own page tables, we can configure different processes to use different microarchitecture oblivious bits and store these bits in the per-process page tables.

B. Microarchitecture Changes

One of the central ideas in *Oreo* is to limit the usage of secret-dependent randomized virtual addresses in microarchitecture structures. We modify the processor pipeline to extensively use masked addresses except for the virtual address security check at the commit stage. We divide the pipeline into three components: the frontend for fetch and decode, the middle component for execute and memory operations, and the backend for committing instructions. In Figure 7, we use different colors to indicate the usage of different types of addresses in each microarchitecture structure: red for secret bits and secret-dependent virtual addresses, green for

secret-free masked addresses, and blue for secret-free physical addresses.

The Fetch & Decode Stage. The frontend fetch stage maintains a PC (program counter) register and the branch predictor. *Oreo* uses the masked address in both the PC register and the branch predictor, requiring no changes to the fetch stage. To understand why, consider the four sources that the hardware uses to update the PC register: (1) PC + instruction size for non-branch instruction; (2) PC + an immediate value from the decode stage for relative direct branches such as `jmp short`; (3) target address predicted by the branch predictor; and (4) target address from the register file or memory, usually for indirect jumps, such as `call` and `ret`.

Among the four sources above, only the last type obtains the branch target from other pipeline stages, namely the execute & memory stage. We place a `Virt2Mask` module between the execute & memory and the fetch stages to ensure the external PC update source uses masked addresses. As such, all the internal updates of the fetch stage, including the PC-derived targets and predicted targets, will consistently use masked addresses without extra intervention.

The Execute & Memory Stage. The middle component of a speculative processor uses a ROB to track all the in-flight instructions, and a load/store queue (LSQ) to track all the in-flight load and store operations.

We extend the ROB (reorder buffer) to store the correct microarchitecture oblivious bits (i.e., $\text{offset}_{\text{oreo}}$) for the PC of each instruction to facilitate *Oreo*’s security check at the commit stage. Specifically, when fetching instructions using a masked address, the address translation procedure looks up TLBs or performs a page table walk to obtain its physical address. As we store $\text{offset}_{\text{oreo}}$ in PTE entries and the TLB, the translation procedure can obtain $\text{offset}_{\text{oreo}}$ and send it back to the core.

In this pipeline stage, we also prevent information leakage caused by load and store instructions. When inserting each load and store instruction into the load/store queue (LSQ), we convert the virtual address used by the load/store instruction into a secret-free masked address and also extract the microarchitecture oblivious bits. Memory dependency checks and other microarchitectural optimizations, such as load-to-store forwarding and address prediction, use masked addresses as inputs. When a load/store is issued to the memory system, similar to instruction fetch, the address translation procedure uses the masked address. For each load/store instruction, the memory system returns data (if it is a load) and the correct $\text{offset}_{\text{oreo}}$ for the given masked address. This correct $\text{offset}_{\text{oreo}}$ is then compared against the extracted microarchitecture oblivious bits in the LSQ entry to determine whether the original load/store virtual address is valid. The comparison result is stored in the corresponding LSQ entry. Note that, this operation is secure because this pre-computed check result does not affect any other microarchitecture states, will only be used in the commit stage, and thus does not introduce new timing side channels.

The Commit Stage. In the baseline hardware, security checks

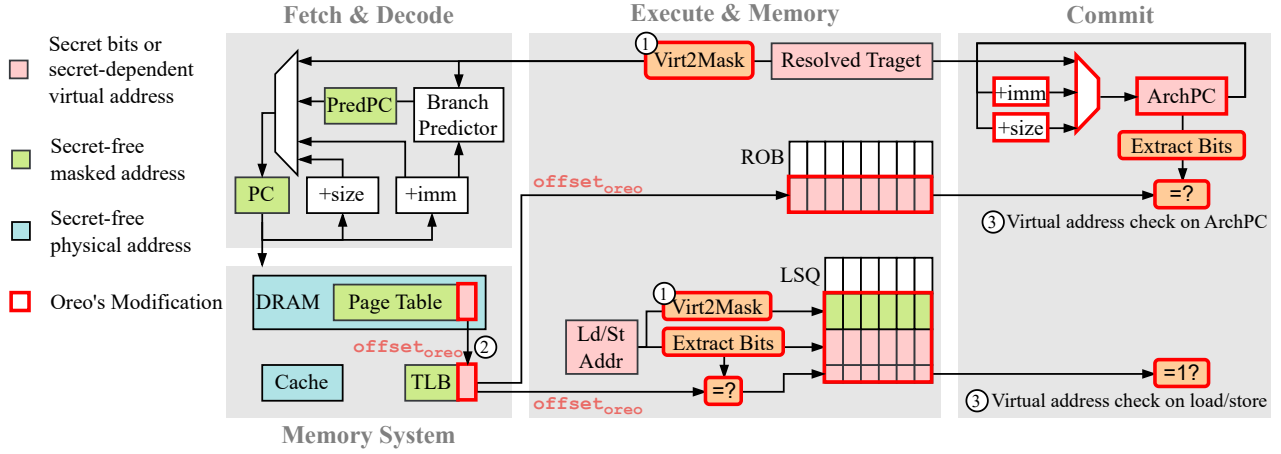


Figure 7: Microarchitecture changes required by *Oreo*.

on memory accesses are performed during the address translation within the MMU. When an instruction enters the commit stage, the hardware checks whether it should trigger an exception (e.g., failing to pass a security check) and retires the instruction if no exception occurs. In *Oreo*, MMU only performs page permission checks on masked addresses during speculation, i.e., whether the masked address has the correct read/write/execute permission on a given page. Notably, the MMU leaves the check of whether the microarchitecture oblivious bits of the virtual address is correct to the commit stage.

We perform two virtual address checks in the commit stage depending on the instruction types. For any type of instruction, we check whether the microarchitecture oblivious bits of its PC is correct. For load/store instructions, we additionally perform this check on the load/store addresses. As we have pre-computed the check result on load/store addresses in the execute & memory stage and stored the check result in the LSQ, we can directly use this check result in the commit stage.

To check the microarchitecture oblivious bits in PCs, we add another PC register to the commit stage. This PC register keeps track of the virtual address for the instruction at the head of ROB. To distinguish from the PC register in the fetch stage, we call it ArchPC, since the PC register at the fetch stage is *speculative*, while the one at the commit stage keeps track of the PC to be committed, reflecting architecture-level information. Since we inserted the correct $offset_{oreo}$ into the ROB in previous stages, we can conveniently check the validity of ArchPC by extracting microarchitecture oblivious bits of the ArchPC and comparing them with the correct bits. If the ArchPC is valid, we can retire the instruction. Otherwise, an exception is triggered for a virtual address check failure. Note that we always let the other exceptions take priority over this virtual address check results to avoid leaking ASLR secrets to attacks using microarchitectural replay attacks [70].

Furthermore, the commit stage replicates some of the next-PC computation logic similar to the fetch stage. As shown

in Figure 7, this replicated part involves integer addition operators, data forwarding from the execution stage, and a multiplexer. After successfully retiring an instruction, we use the next-PC logic to derive the architecture PC value of the next instruction and update the ArchPC register. The design above is a naive baseline to showcase that *Oreo*'s hardware modification is not intrusive. There exist plenty of optimization opportunities to omit these replicated computation structures or move it to an earlier stage of the pipeline.

C. Timing and Hardware Cost Analysis

We provide a comprehensive analysis to estimate the timing and area overhead. We acknowledge that a synthesizable implementation of the hardware components of *Oreo* can offer a more accurate measurement.

Pipeline Timing Impacts. The hardware changes introduced by *Oreo* include: (1) applying *Virt2Mask* to virtual addresses to obtain masked addresses; (2) obtaining $offset_{oreo}$ from page tables/TLBs; (3) checking whether a virtual address has valid microarchitecture oblivious bits at the commit stage. We labeled these changes in Figure 7.

First, the *Virt2Mask* operation is performed before branch squashing and issuing load/store requests, labeled as ① in Figure 7. Given there can exist multiple randomization regions, when converting a virtual address to a masked address, we first determine which randomization region the input virtual address falls into. This operation requires parallel comparisons between the virtual address and the boundaries of each randomization region. Once the randomization region is determined, we obtain the information of which bits are microarchitecture oblivious bits. We can then clear these bits using a quick bit-wise AND operation to obtain the masked address. Overall, it is a lightweight operation consisting of parallel comparisons and bit-wise AND operations. Hence, we count no extra cycles introduced by *Virt2Mask*.

Second, *Oreo* obtains secret ASLR offsets from page tables to fill in TLB entries and further passes them to the ROB and

the load/store unit, labeled as ② in Figure 7. These operations are performed in parallel with address translation and do not introduce extra latency.

Third, *Oreo* extracts microarchitecture oblivious bits in the ArchPC register and checks whether they match the correct bits or not, labeled as ③ in Figure 7. We count no extra latency or back-pressure at commit time because the information used for the virtual address check on the instruction PC is available at the execution stage, and the logic to do the check is a simple combinational circuit. Furthermore, the virtual address check results for load/store addresses are pre-computed before commit time as discussed in Section VI-B.

Area Cost. We analyze the area cost of *Oreo*'s hardware by visiting each block in Figure 7. To begin with, no extra hardware is needed in the fetch stage as this stage purely operates on masked addresses.

In the memory system, *Oreo* includes additional bits to record the correct $\text{offset}_{\text{oreo}}$ for each TLB entry. Specifically, in our prototyped system, we add 8 extra bits per TLB entry. For reference, a Mega BOOM processor [83] has 584 TLB entries (including iTLB, dTLB, and L2TLB), leading to an overhead of 584 bytes.

In the execute and memory stage of the pipeline, *Oreo* incorporates the Virt2Mask and bits extraction modules, as well as extra bits in the ROB and LSQ. First, both the Virt2Mask module and the bits extraction module need to store metadata of randomization regions, including the boundary and microarchitecture oblivious bits information. For each randomization region, we use 128 bits to store its boundary (i.e., start and end), and a 64-bit vector to indicate which bits are microarchitecture oblivious bits. Our prototype implementation uses two randomization regions to protect the kernel and user ASLR, which yields a total 384-bit overhead. Second, *Oreo* introduces extra fields in the ROB and LSQ to facilitate virtual address checks on instruction PCs and load/store addresses, including 8 bits to store the correct $\text{offset}_{\text{oreo}}$ for PC in each ROB entry, another 8 bits to store the extracted offset bits of the load/store address in each LSQ entry, and 1 bit per LSQ entry to store the pre-computed virtual address check result. For reference, the Mega BOOM processor has a 128-entry ROB and a 64-entry LSQ, resulting in 200 bytes storage overhead.

In the commit stage, we add the ArchPC register, which is 64-bit. The bits extraction module requires extra storage for recording randomization region metadata, which is shared with the modules in the execute and memory stages.

Overall, using the Mega BOOM processor as an example, *Oreo* incurs small storage overhead, including 256 bytes in-core overhead and 584 bytes overhead in the memory system.

VII. EVALUATION

A. Experiment Setup

We implement our kernel changes in Linux 6.6. We use a kernel patch by Hou et. al. [32] to relocate kernel text and modules to a 444GB unused region in the kernel address space. We implement our microarchitecture changes in the




Figure 8: Performance evaluation results on SPEC2017

gem5 simulator (v24.0) [9], [54] using the full system mode. The microarchitecture configuration is similar to the configurations used in prior microarchitectural mitigation papers [53], [77], [82]. We model a 1-core CPU for running the SPEC2017 benchmark [11] and security evaluation, and a 2-core CPU for running the LEBench benchmark [85]. We configure each core as an 8-issue out-of-order (O3) superscalar processor with 32 load queue entries, 32 store queue entries, and 192 ROB entries. The branch predictor uses the tournament prediction policy with 4096 BTB entries and 16 RAS entries. We model 64KB 8-way L1 I-cache and D-cache and a 2MB 16-way L2 cache. gem5 has a customized procedure for booting Linux which deviates from how Linux boots on a real processor and does not support kernel ASLR. In our implementation, we modify gem5's booting procedure to support kernel ASLR. Overall, our prototype involves 785 lines of code (LoC) changes to the Linux kernel, and 1897 LoC modifications to the gem5 simulator.

B. Performance Results

SPEC2017. We evaluate the performance overhead of *Oreo* on the SPEC2017 IntRate benchmark [11]. We configure the applications to use reference input size, warm up the system and microarchitecture structures by executing 10 billion user-space instructions, and then measure the performance of the next 1 billion instructions. We skipped 557.xz_r because the simulation crashes due to a bug in gem5 for not supporting certain instructions used by this application. Figure 8 shows the reported CPI (cycles per user-space instruction) for each application. We label the CPI overhead ratios incurred by *Oreo* compared to the baseline on top of the green bars. The overall CPI (counting both user-space and kernel instructions) is mostly identical to the CPI for userspace only, given that the userspace time dominates when running the SPEC benchmark.

Across all applications, *Oreo* introduces negligible performance overhead compared to the baseline, incurring 0.11% CPI overhead on average. This indicates that *Oreo*'s changes on the software and hardware have little impact on the overall performance of user-space applications.

We report the ratio of memory accesses where *Oreo* applies its protection to provide insight into how much of the program execution triggers protection. This information allows us to validate the relevance of the SPEC benchmark in evaluating the performance impacts of *Oreo*. In our prototype, we apply Virt2Mask to any addresses that fall within the user address




Figure 9: Performance evaluation results on LEBench

space, or the selected randomization region for kernel code or modules. For the SPEC benchmark, at least 99.46% of the memory accesses trigger *Oreo*'s protection, with this being the minimum ratio observed among all applications.

LEBench. We also evaluate the performance overhead of *Oreo* using the LEBench benchmark [85], a microbenchmark suite that measures the performance of kernel operations (system calls). We run the LEBench measurement for each system call 16 times and report the average latency overhead in the top half of Figure 9, comparing *Oreo* and the insecure baseline. *Oreo* introduces an average overhead -0.28% across all system calls in the suite, which is almost negligible.

However, unlike the SPEC benchmark, we observe large performance variations in the LEBench, ranging from -6.39% to 6.16% . To understand the variation, for each system call, we plot the range of multiple measured latencies normalized to the medium latency of the insecure baseline in the bottom half of Figure 9. Specifically, the colored box indicates the first quartile (Q1) and the third quartile (Q3) of normalized latency, with the medium latency marked as a horizontal line dividing the box into two halves. We additionally use hollow dots to mark outliers.

From the figure, we observe that large variations consistently exist for certain system calls, such as `thrcrcreate` and `pagefault`, in the baseline and when using *Oreo*. Given that these system calls require coordination between multiple threads, we suspect that the large variation is caused by the dynamic non-deterministic scheduling of the Linux kernel.

Another difference between LEBench and SPEC2017 is the ratio of memory accesses that use *Oreo*'s protection. In the LEBench benchmark, the average ratio of memory accesses that need `Virt2Mask` is 75.17%, ranging from 65.97% to 82.86%, which is lower than the ratio in SPEC2017. This is because LEBench contains more memory accesses to kernel data, which is not protected by our prototype implementation.




Figure 10: Evaluating the prefetch attack on the insecure baseline and *Oreo*.

C. Security Evaluation

We conduct three experiments to validate the security properties of *Oreo* and to demonstrate that our implementation aligns with the design presented in the paper. Additionally, we provide a formal proof in a technical report [71] to show that *Oreo* achieves a non-interference property, preventing the leakage of ASLR secrets.

The Prefetch Attack. We evaluate a prefetch attack [31], [49], [52] on the insecure baseline and *Oreo*. In both cases, the kernel code (including text and modules) is randomly relocated to the range from `0xfffffff8000000000` to `0xfffffffef000000000`. The attacker scans the randomization region by probing addresses with a stride of 2 GB. The probing operation executes the prefetch instruction twice. The first fetch operation brings the address into various microarchitecture structures. The attacker then measures the latency of the second fetch operation.

Figure 10 shows the attack results. The prefetch attack works effectively on the insecure baseline, where the prefetch latency is distinctively lower at the randomized kernel address `0xfffffff8601800040`, and is consistently high at the other unmapped kernel addresses. This timing difference is caused by the fact that TLB caches address translation only for mapped addresses, not for unmapped addresses. In contrast, the prefetch attack no longer works on *Oreo*. *Oreo* converts virtual addresses to masked addresses and uses masked addresses to access the TLB. Since the TLB already cached the corresponding masked address during the execution of the first prefetch operation, the second prefetch always results in a hit, resulting in indistinguishable low latency.

Leakage Path ①. In addition to the prefetch attack, we evaluate the speculative code region probing attack from the BlindSide paper [27]. We consider this attack as a representative attack for the leakage path ①. Besides, we think solely relying on timing may not fully capture the effectiveness of our design. Thus, we validate the attack results by comparing internal microarchitectural traces dumped from the `gem5` simulator.

Figure 11(a) shows the attack gadget. Following the attack described in BlindSide [27], we assume an attacker leverages a memory corruption vulnerability to overwrite a function pointer in the kernel memory. The attacker then triggers the victim to transiently jump to the corrupted function pointer during the mis-speculation of a conditional branch. In our experiment, we set the function pointer to a valid address and an invalid address and compare the input traces on TLB, cache,

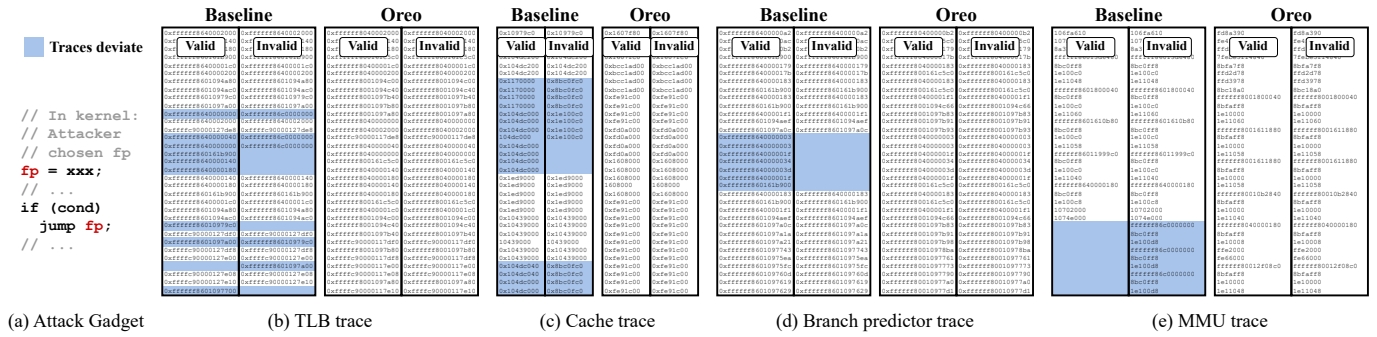


Figure 11: Evaluating the speculative code gadget probing attack on the insecure baseline and *Oreo*. (b)-(e) compare the input traces of multiple microarchitecture structures when transiently jumping to a valid and an invalid address.

branch predictor, and MMU, shown in Figure 11(b)-(e). We highlight the parts where the two traces deviate from each other, indicating an exploitable side channel.

On the baseline, the input traces to the four microarchitecture structures all deviate when transiently accessing the valid and the invalid addresses. In contrast, with *Oreo*, the traces are identical since the same masked address is used for transient access. This indicates the attacker cannot distinguish valid addresses from invalid ones through transient memory accesses, so *Oreo* can successfully block leakage path ①.

Leakage Path ②. We use microarchitectural traces, similar to Figure 11, to evaluate *Oreo* against the second leakage path. We trigger the victim kernel to execute a system call and record the input traces to various microarchitecture structures. We then examine these traces and find that on the baseline, the input traces to these microarchitecture structures all include addresses with secret $\text{offset}_{\text{oreo}}$. With *Oreo*, the input traces only consist of addresses without these secret bits.

VIII. RELATED WORK

We discuss related work aimed at strengthening the security property of ASLR schemes. We first discuss ASLR protection schemes that aim to address different information leakage threats, i.e., through microarchitectural attacks or software-level attacks. We then discuss mechanisms designed to increase ASLR entropy.

Blocking Software-Based ASLR Bypasses. Several prior work [7], [17], [18], [26], [55], [73], [74] aim to make it more difficult for attackers to leak ASLR secrets via exploiting software vulnerabilities. For example, ASLR-Guard [55] uses encryption to prevent leaking code pointers. KASLR-MT [74] and Vano-Garcia et al. [73] block information leakage due to memory deduplication attacks. XnR [7] and Readactor [17], [18], and Gionta et al. [26] enable “executing-only-memory” to prevent reading and then leaking code pointers. These techniques focus on software-level threats and are ineffective against microarchitectural attacks. They can complement *Oreo* to further strengthen the security of ASLR schemes.

Blocking Microarchitectural-Attack-Assisted ASLR Bypasses. This group of mitigation mechanisms [13], [25], [30] shares the same goal as our work. Several defenses

aim to prevent ASLR secrets from being leaked via virtual memory layout probing attacks, blocking leakage path ①. One approach is to isolate the kernel and user-space address spaces to prevent memory layout probing, as demonstrated in KAISER [30] (also known as KPTI) and LAZARUS [25]. However, given that some of the kernel trampoline pages are still mapped in the user space, the Linux prototype of KAISER [1] is still vulnerable to ASLR bypasses, as shown in EchoLoad [30] and Entrybleed [52]. Besides, software-level isolation does not help mitigate the second leakage path, where the ASLR randomized bits are leaked when the victim program uses secret pointers as program counters or load/store addresses.

Alternatively, FLARE [13] makes accessing valid (mapped) and invalid (unmapped) kernel addresses take the same amount of time. It works by mapping all invalid kernel addresses to one valid physical page, so that accessing these invalid pages will need to go through the full page table walk. This mitigation has several limitations. First, it cannot block memory layout probing attacks that use BTB and TLB as side channels. Second, given that the invalid addresses now map to an empty physical page, whose content is different from the actual valid pages, any speculative data-dependent accesses can be used to distinguish between invalid and valid addresses. In contrast to this ad-hoc solution, *Oreo* takes a much more comprehensive approach to make accessing invalid and valid addresses exhibit indistinguishable side effects as long as they map to the same masked address.

Increasing ASLR Entropy. Several prior works aim to increase the entropy of ASLR [14], [17], [24], [46], [59] by introducing finer-grained ASLR, increasing the size of the randomization region, and re-randomization. For example, FGKASLR [46] randomizes function orders in kernel code to make it harder to find code gadgets. Adelle [59] increases the kernel ASLR randomization region to the entire 64-bit virtual memory and re-randomizes the layout of kernel modules. Readactor [17], [18] and CodeArmor [14] re-randomizes code pointers or code memory layout. Morpheus [24] periodically re-randomizes the code pointers with a higher frequency than previous approaches as it leverages complex hardware modifications to do the re-randomization.

Power-Induced Timing Side Channels. Recent work [49] breaks ASLR through power-induced timing side channels. Although it is out of *Oreo*'s protection scope, *Oreo*'s implementation can make it easier to mitigate this class of side channels. Power-based side channels arise when circuits perform computation using secret-dependent values. Different inputs activate different transistors and lead to different amounts of power consumption. According to Figure 7, the protected secret bits are extracted from virtual addresses or obtained from TLB entries, and are only used in two equality checks (denoted by the “=?” boxes). Thus, it becomes feasible to adopt classic mitigations, such as circuit masking [57], to secure these specific modules against power side channels.

IX. CONCLUSION

This paper presented a systematic analysis of microarchitectural side-channel attacks for ASLR bypasses. We use our analysis to guide the design of *Oreo*. *Oreo* strengthens the security of ASLR via the modification to the memory interface. By introducing an extra layer of masked address space and converting virtual addresses to masked addresses to be used to set up page tables and index into various microarchitecture structures, *Oreo* constrains the secret offset exposure in both software and hardware. Our security and performance evaluation demonstrates that *Oreo* is an effective mitigation and introduces negligible performance overhead.

ACKNOWLEDGMENT

The authors thank the Matcha Group (MIT) for their help and the anonymous NDSS reviewers for their feedback. This work was supported in part by a gift from Amazon; by the Air Force Office of Scientific Research (AFOSR) under grant FA9550-22-1-0511; by ACE, one of the seven centers in JUMP 2.0, a Semiconductor Research Corporation (SRC) program sponsored by DARPA.

REFERENCES

- [1] “Page table isolation (PTI),” <https://www.kernel.org/doc/html/next/x86/pti.html>.
- [2] “2023 CWE top 25 most dangerous software weaknesses,” https://cwe.mitre.org/top25/archive/2023/2023_top25_list.html, 2023.
- [3] “Memory management,” https://www.kernel.org/doc/html/v6.4/arch/x86/x86_64/mm.html, 2024.
- [4] “Memory protection keys,” <https://docs.kernel.org/core-api/protection-keys.html>, 2024.
- [5] S. Ainsworth, “GhostMinion: A strictness-ordered cache system for spectre mitigation,” in *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture*, 2021, pp. 592–606.
- [6] S. Ainsworth and T. M. Jones, “MuonTrap: Preventing cross-domain spectre-like attacks by capturing speculative state,” in *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 2020, pp. 132–144.
- [7] M. Backes, T. Holz, B. Kollenda, P. Koppe, S. Nürnberg, and J. Pewny, “You can run but you can’t read: Preventing disclosure exploits in executable code,” in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, 2014, pp. 1342–1353.
- [8] K. Barber, A. Bacha, L. Zhou, Y. Zhang, and R. Teodorescu, “SpecShield: Shielding speculative data from microarchitectural covert channels,” in *2019 28th International Conference on Parallel Architectures and Compilation Techniques (PACT)*. IEEE, 2019, pp. 151–164.
- [9] N. L. Binkert, B. M. Beckmann, G. Black, S. K. Reinhardt, A. G. Saida, A. Basu, J. Hestness, D. Hower, T. Krishna, S. Sardashti, R. Sen, K. Sewell, M. S. B. Altaf, N. Vaish, M. D. Hill, and D. A. Wood, “The gem5 simulator,” *SIGARCH Comput. Archit. News*, vol. 39, no. 2, pp. 1–7, 2011. [Online]. Available: <https://doi.org/10.1145/2024716.2024718>
- [10] T. Bletsch, X. Jiang, V. W. Freeh, and Z. Liang, “Jump-oriented programming: a new class of code-reuse attack,” in *Proceedings of the 6th ACM symposium on information, computer and communications security*, 2011, pp. 30–40.
- [11] J. Bucek, K.-D. Lange, and J. v. Kistowski, “SPEC CPU2017: Next-generation compute benchmark,” in *Companion of the 2018 ACM/SPEC International Conference on Performance Engineering*, 2018, pp. 41–42.
- [12] C. Canella, D. Genkin, L. Giner, D. Gruss, M. Lipp, M. Minkin, D. Moghimi, F. Piessens, M. Schwarz, B. Sunar *et al.*, “Fallout: Leaking data on Meltdown-resistant cpus,” in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 2019, pp. 769–784.
- [13] C. Canella, M. Schwarz, M. Haubenwallner, M. Schwarzl, and D. Gruss, “KASLR: Break it, fix it, repeat,” in *Proceedings of the 15th ACM Asia Conference on Computer and Communications Security*, 2020, pp. 481–493.
- [14] X. Chen, H. Bos, and C. Giuffrida, “CodeArmor: Virtualizing the code space to counter disclosure attacks,” in *2017 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2017, pp. 514–529.
- [15] R. Choudhary, J. Yu, C. Fletcher, and A. Morrison, “Speculative privacy tracking (SPT): Leaking information from speculative execution without compromising privacy,” in *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture*, 2021, pp. 607–622.
- [16] M. H. I. Chowdhury and F. Yao, “Leaking secrets through modern branch predictors in the speculative world,” *IEEE Transactions on Computers*, vol. 71, no. 9, pp. 2059–2072, 2021.
- [17] S. Crane, C. Liebchen, A. Homescu, L. Davi, P. Larsen, A.-R. Sadeghi, S. Brunthaler, and M. Franz, “Readactor: Practical code randomization resilient to memory disclosure,” in *2015 IEEE Symposium on Security and Privacy*. IEEE, 2015, pp. 763–780.
- [18] S. J. Crane, S. Volckaert, F. Schuster, C. Liebchen, P. Larsen, L. Davi, A.-R. Sadeghi, T. Holz, B. De Sutter, and M. Franz, “It’s a TRaP: Table randomization and protection against function-reuse attacks,” in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, 2015, pp. 243–255.
- [19] M. Dai, R. Paccagnella, M. Gomez-Garcia, J. McCalpin, and M. Yan, “Don’t mesh around: Side-Channel attacks and mitigations on mesh interconnects,” in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 2857–2874.
- [20] L.-A. Daniel, M. Bogner, J. Noorman, S. Bardin, T. Rezk, and F. Piessens, “ProSpeCT: Provably secure speculation for the Constant-Time policy,” in *32nd USENIX Security Symposium (USENIX Security 23)*, 2023, pp. 7161–7178.
- [21] D. Evtvushkin, D. Ponomarev, and N. Abu-Ghazaleh, “Jump over ASLR: Attacking branch predictors to bypass ASLR,” in *2016 49th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 2016, pp. 1–13.
- [22] D. Evtvushkin, R. Riley, N. C. Abu-Ghazaleh, ECE, and D. Ponomarev, “BranchScope: A new side-channel attack on directional branch predictor,” *ACM SIGPLAN Notices*, vol. 53, no. 2, pp. 693–707, 2018.
- [23] J. Fustos, F. Farshchi, and H. Yun, “SpectreGuard: An efficient data-centric defense mechanism against Spectre attacks,” in *Proceedings of the 56th Annual Design Automation Conference 2019*, 2019, pp. 1–6.
- [24] M. Gallagher, L. Biernacki, S. Chen, Z. B. Aweke, S. F. Yitbarek, M. T. Aga, A. Harris, Z. Xu, B. Kasikci, V. Bertacco *et al.*, “Morpheus: A vulnerability-tolerant secure architecture based on ensembles of moving target defenses with churn,” in *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, 2019, pp. 469–484.
- [25] D. Gens, O. Arias, D. Sullivan, C. Liebchen, Y. Jin, and A.-R. Sadeghi, “Lazarus: Practical side-channel resilient kernel-space randomization,” in *Research in Attacks, Intrusions, and Defenses: 20th International Symposium, RAID 2017, Atlanta, GA, USA, September 18–20, 2017, Proceedings*. Springer, 2017, pp. 238–258.
- [26] J. Gionta, W. Enck, and P. Larsen, “Preventing kernel code-reuse attacks through disclosure resistant code diversification,” in *2016 IEEE Conference on Communications and Network Security (CNS)*. IEEE, 2016, pp. 189–197.

- [27] E. Göktaş, K. Razavi, G. Portokalidis, H. Bos, and C. Giuffrida, “Speculative probing: Hacking blind in the spectre era,” in *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, 2020, pp. 1871–1885.
- [28] B. Gras, K. Razavi, H. Bos, and C. Giuffrida, “Translation leak-aside buffer: Defeating cache side-channel protections with TLB attacks,” in *27th USENIX Security Symposium (USENIX Security 18)*, 2018, pp. 955–972.
- [29] B. Gras, K. Razavi, E. Bosman, H. Bos, and C. Giuffrida, “ASLR on the line: Practical cache attacks on the MMU,” in *NDSS*, vol. 17, 2017, p. 26.
- [30] D. Gruss, M. Lipp, M. Schwarz, R. Fellner, C. Maurice, and S. Mangard, “KASLR is dead: Long live KASLR,” in *Engineering Secure Software and Systems: 9th International Symposium, ESSoS 2017, Bonn, Germany, July 3-5, 2017, Proceedings 9*. Springer, 2017, pp. 161–176.
- [31] D. Gruss, C. Maurice, A. Fogh, M. Lipp, and S. Mangard, “Prefetch side-channel attacks: Bypassing SMAP and kernel ASLR,” in *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 2016, pp. 368–379.
- [32] W. Hou, T. Garnier, and B. Gerst, “x86/pie: Make kernel image’s virtual address flexible,” <https://lore.kernel.org/lkml/cover.1682673542.git.houwenlong.hwl@antgroup.com/>, 2023.
- [33] R. Hund, C. Willems, and T. Holz, “Practical timing side channel attacks against kernel space ASLR,” in *2013 IEEE Symposium on Security and Privacy*. IEEE, 2013, pp. 191–205.
- [34] Inc. Qualcomm Technologies, “Pointer authentication on ARMv8.3: Design and analysis of the new software security instructions,” <https://www.qualcomm.com/content/dam/qcomm-martech/dm-assets/documents/pointer-auth-v7.pdf>, 2017.
- [35] C. Intel, “Intel® 64 and IA-32 architectures software developer’s manual volume 3 (3a, 3b, 3c & 3d): System programming guide,” *Denver, [2006]*, 2022.
- [36] Jake Edge, “Kernel address space layout randomization,” <https://lwn.net/Articles/569635/>, 2013.
- [37] Y. Jang, S. Lee, and T. Kim, “Breaking kernel address space layout randomization with Intel TSX,” in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 2016, pp. 380–392.
- [38] S. Jenkins, “Exploiting CVE-2022-42703 - bringing back the stack attack,” <https://googleprojectzero.blogspot.com/2022/12/exploiting-CVE-2022-42703-bringing-back-the-stack-attack.html>, 2022.
- [39] K. N. Khasawneh, E. M. Koruyeh, C. Song, D. Evtvushkin, D. Ponomarev, and N. Abu-Ghazaleh, “SafeSpec: Banishing the Spectre of a Meltdown with leakage-free speculation,” in *2019 56th ACM/IEEE Design Automation Conference (DAC)*. IEEE, 2019, pp. 1–6.
- [40] V. Kiriansky, I. Lebedev, S. Amarasinghe, S. Devadas, and J. Emer, “DAWG: A defense against cache timing attacks in speculative execution processors,” in *2018 51st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 2018, pp. 974–987.
- [41] V. Kiriansky and C. Waldspurger, “Speculative buffer overflows: Attacks and defenses,” *arXiv preprint arXiv:1807.03757*, 2018.
- [42] P. Kocher, J. Horn, A. Fogh, D. Genkin, D. Gruss, W. Haas, M. Hamburg, M. Lipp, S. Mangard, T. Prescher, M. Schwarz, and Y. Yarom, “Spectre attacks: Exploiting speculative execution,” in *40th IEEE Symposium on Security and Privacy (S&P’19)*, 2019.
- [43] E. M. Koruyeh, K. N. Khasawneh, C. Song, and N. Abu-Ghazaleh, “Spectre returns! speculation attacks using the return stack buffer,” in *12th USENIX Workshop on Offensive Technologies (WOOT 18)*, 2018.
- [44] E. M. Koruyeh, S. H. A. Shirazi, K. N. Khasawneh, C. Song, and N. Abu-Ghazaleh, “SpecCFI: Mitigating Spectre attacks using CFI informed speculation,” in *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2020, pp. 39–53.
- [45] J. Koschel, C. Giuffrida, H. Bos, and K. Razavi, “TagBleed: Breaking KASLR on the isolated kernel address space using tagged TLBs,” in *2020 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2020, pp. 309–321.
- [46] Kristen Carlson Accardi, “Finer grained kernel address space randomization,” <https://lwn.net/Articles/811685/>, 2020.
- [47] S. Lee, M.-W. Shih, P. Gera, T. Kim, H. Kim, and M. Peinado, “Inferring fine-grained control flow inside SGX enclaves with branch shadowing,” in *26th USENIX Security Symposium (USENIX Security 17)*, 2017, pp. 557–574.
- [48] P. Li, L. Zhao, R. Hou, L. Zhang, and D. Meng, “Conditional speculation: An effective approach to safeguard out-of-order execution against Spectre attacks,” in *2019 IEEE international symposium on high performance computer architecture (HPCA)*. IEEE, 2019, pp. 264–276.
- [49] M. Lipp, D. Gruss, and M. Schwarz, “AMD prefetch attacks through power and time,” in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 643–660.
- [50] M. Lipp, V. Hadžić, M. Schwarz, A. Perais, C. Maurice, and D. Gruss, “Take a way: Exploring the security implications of AMD’s cache way predictors,” in *Proceedings of the 15th ACM Asia Conference on Computer and Communications Security*, 2020, pp. 813–825.
- [51] M. Lipp, M. Schwarz, D. Gruss, T. Prescher, W. Haas, A. Fogh, J. Horn, S. Mangard, P. Kocher, D. Genkin, Y. Yarom, and M. Hamburg, “Meltdown: Reading kernel memory from user space,” in *27th USENIX Security Symposium (USENIX Security 18)*, 2018.
- [52] W. Liu, J. Ravichandran, and M. Yan, “EntryBleed: A universal KASLR bypass against KPTI on Linux,” in *Proceedings of the 12th International Workshop on Hardware and Architectural Support for Security and Privacy*, 2023, pp. 10–18.
- [53] K. Loughlin, I. Neal, J. Ma, E. Tsai, O. Weisse, S. Narayanasamy, and B. Kasikci, “DOLMA: Securing speculation with the principle of transient Non-Observability,” in *30th USENIX Security Symposium (USENIX Security 21)*, 2021, pp. 1397–1414.
- [54] J. Lowe-Power, A. M. Ahmad, A. Akram, M. Alian, R. Amslinger, M. Andreozzi, A. Armejach, N. Asmussen, S. Bharadwaj, G. Black, G. Bloom, B. R. Bruce, D. R. Carvalho, J. Castrill’on, L. Chen, N. Derumigny, S. Diestelhorst, W. Elsasser, M. Fariborz, A. F. Farahani, P. Fotouhi, R. Gambord, J. Gandhi, D. Gope, T. Grass, B. Hanindhito, A. Hansson, S. Haria, A. Harris, T. Hayes, A. Herrera, M. Horsnell, S. A. R. Jafri, R. Jagtap, H. Jang, R. Jeyapaul, T. M. Jones, M. Jung, S. Kannoth, H. Khaleghzadeh, Y. Kodama, T. Krishna, T. Marinelli, C. Menard, A. Mondelli, T. M’uck, O. Naji, K. Nathella, H. Nguyen, N. Nikoleris, L. E. Olson, M. S. Orr, B. Pham, P. Prieto, T. Reddy, A. Roelke, M. Samani, A. Sandberg, J. Setoain, B. Shingarov, M. D. Sinclair, T. Ta, R. Thakur, G. Travaglini, M. Upton, N. Vaish, I. Vougioukas, Z. Wang, N. Wehn, C. Weis, D. A. Wood, H. Yoon, and Eder F. Zulian, “The gem5 simulator: Version 20.0+,” *CoRR*, vol. abs/2007.03152, 2020. [Online]. Available: <https://arxiv.org/abs/2007.03152>
- [55] K. Lu, C. Song, B. Lee, S. P. Chung, T. Kim, and W. Lee, “ASLR-Guard: Stopping address space leakage for code reuse attacks,” in *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*, 2015, pp. 280–291.
- [56] G. Maisuradze and C. Rossow, “ret2spec: Speculative execution using return stack buffers,” in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, 2018, pp. 2109–2122.
- [57] S. Mangard, E. Oswald, and T. Popp, *Power Analysis Attacks: Revealing the Secrets of Smart Cards*. Springer Science & Business Media, 2008, vol. 31.
- [58] N. Mosier, H. Nemati, J. C. Mitchell, and C. Trippel, “Serberus: Protecting cryptographic code from Spectres at compile-time,” in *2024 IEEE Symposium on Security and Privacy (SP)*. Los Alamitos, CA, USA: IEEE Computer Society, may 2024, pp. 51–51. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/SP54263.2024.00048>
- [59] R. Nikolaev, H. Nadeem, C. Stone, and B. Ravindran, “Adelie: continuous address space layout re-randomization for Linux drivers,” in *Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, 2022, pp. 483–498.
- [60] PaX team, “PaX address space layout randomization (ASLR).” <http://pax.grsecurity.net/docs/aslr.txt>, 2003.
- [61] J. Ravichandran and M. Wang, “Lord of the io_urings,” <https://css.csail.mit.edu/6.5660/2023/projects/jravi-mi27950.pdf>, Tech. Rep., 2022.
- [62] R. Roemer, E. Buchanan, H. Shacham, and S. Savage, “Return-oriented programming: Systems, languages, and applications,” *ACM Transactions on Information and System Security (TISSEC)*, vol. 15, no. 1, pp. 1–34, 2012.
- [63] G. Saileshwar and M. K. Qureshi, “CleanupSpec: An “undo” approach to safe speculation,” in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, 2019, pp. 73–86.
- [64] C. Sakalis, M. Alipour, A. Ros, A. Jimborean, S. Kaxiras, and M. Sjölander, “Ghost loads: What is the cost of invisible speculation?” in

- Proceedings of the 16th ACM International Conference on Computing Frontiers*, 2019, pp. 153–163.
- [65] C. Sakalis, S. Kaxiras, A. Ros, A. Jimborean, and M. Sjölander, “Efficient invisible speculative execution through selective delay and value prediction,” in *Proceedings of the 46th International Symposium on Computer Architecture*, 2019, pp. 723–735.
- [66] M. Schwarz, C. Canella, L. Giner, and D. Gruss, “Store-to-Leak Forwarding: Leaking data on Meltdown-resistant cpus (updated and extended version),” *arXiv preprint arXiv:1905.05725*, 2019.
- [67] M. Schwarz, M. Lipp, C. A. Canella, R. Schilling, F. Kargl, and D. Gruss, “ConTExT: A generic approach for mitigating Spectre,” in *Network and Distributed System Security Symposium 2020*, 2020.
- [68] M. Schwarz, M. Schwarzl, M. Lipp, J. Masters, and D. Gruss, “Net-Spectre: Read arbitrary memory over network,” in *Computer Security—ESORICS 2019: 24th European Symposium on Research in Computer Security, Luxembourg, September 23–27, 2019, Proceedings, Part 1 24*. Springer, 2019, pp. 279–299.
- [69] Siguza, “IOHIDEous,” <https://blog.siguza.net/IOHIDEous/>, 2017.
- [70] D. Skarlatos, M. Yan, B. Gopireddy, R. Sprabery, J. Torrellas, and C. W. Fletcher, “MicroScope: enabling microarchitectural replay attacks,” in *Proceedings of the 46th International Symposium on Computer Architecture*, 2019, pp. 318–331.
- [71] S. Song, J. Zhang, and M. Yan, “Oreo: Protecting ASLR against microarchitectural attacks (extended version),” 2024. [Online]. Available: <https://arxiv.org/abs/2412.07135>
- [72] L. Szekeeres, M. Payer, T. Wei, and D. Song, “SoK: Eternal war in memory,” in *2013 IEEE Symposium on Security and Privacy*. IEEE, 2013, pp. 48–62.
- [73] F. Vano-Garcia and H. Marco-Gisbert, “An info-leak resistant kernel randomization for virtualized systems,” *IEEE Access*, vol. 8, pp. 161 612–161 629, 2020.
- [74] —, “KASLR-MT: Kernel address space layout randomization for multi-tenant cloud systems,” *Journal of Parallel and Distributed Computing*, vol. 137, pp. 77–90, 2020.
- [75] Y. Wang, R. Paccagnella, E. T. He, H. Shacham, C. W. Fletcher, and D. Kohlbrenner, “Hertzbleed: Turning power Side-Channel attacks into remote timing attacks on x86,” in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 679–697.
- [76] D. Weber, A. Ibrahim, H. Nemati, M. Schwarz, and C. Rossow, “Osiris: Automated discovery of microarchitectural side channels,” in *30th USENIX Security Symposium (USENIX Security 21)*, 2021, pp. 1415–1432.
- [77] O. Weisse, I. Neal, K. Loughlin, T. F. Wenisch, and B. Kasikci, “NDA: Preventing speculative execution attacks at their source,” in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, 2019, pp. 572–586.
- [78] J. Wikner, D. Trujillo, and K. Razavi, “Phantom: Exploiting decoder-detectable mispredictions,” in *MICRO Conference 2023*, 2023.
- [79] M. Yan, J. Choi, D. Skarlatos, A. Morrison, C. Fletcher, and J. Torrellas, “InvisiSpec: Making speculative execution invisible in the cache hierarchy,” in *2018 51st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 2018, pp. 428–441.
- [80] J. Yu, L. Hsiung, M. El Hajj, and C. W. Fletcher, “Data oblivious ISA extensions for side channel-resistant and high performance computing,” *Cryptology ePrint Archive*, 2018.
- [81] J. Yu, N. Mantri, J. Torrellas, A. Morrison, and C. W. Fletcher, “Speculative data-oblivious execution: Mobilizing safe prediction for safe and efficient speculative execution,” in *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 2020, pp. 707–720.
- [82] J. Yu, M. Yan, A. Khyzha, A. Morrison, J. Torrellas, and C. W. Fletcher, “Speculative taint tracking (STT): A comprehensive protection for speculatively accessed data,” in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, 2019, pp. 954–968.
- [83] J. Zhao, B. Korpan, A. Gonzalez, and K. Asanovic, “SonicBOOM: The 3rd generation Berkeley Out-of-Order Machine,” May 2020.
- [84] Z. N. Zhao, A. Morrison, C. W. Fletcher, and J. Torrellas, “Binoculars: Contention-Based Side-Channel attacks exploiting the page walker,” in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 699–716.
- [85] Zirui Zhao, “LEBench-Sim: A Benchmark Suite for Large-Eddy Simulation,” <https://github.com/zzrcxb/LEBench-Sim>, 2022.

A. Description & Requirements

1) *How to access*: The artifact for reproducing the results in the paper is at <https://doi.org/10.5281/zenodo.14261065>. The source code of our implementation and experiments is available at <https://github.com/CSAIL-Arch-Sec/Oreo>.

2) *Hardware dependencies*: The artifact utilizes the gem5 simulator (v24.0) [9], [54] to emulate Linux and execute benchmarks in full-system mode, requiring only CPU, memory, and disk resources. Each gem5 instance needs 1 core and 2 GB memory. The artifact’s resources, including the source files and a generated disk image, require 50 GB of available disk space. We estimate the runtime of each experiment in this artifact on an Intel(R) Xeon(R) Gold 5220R CPU with 252 GB memory.

3) *Software dependencies*: The artifact requires Vagrant (v2.4) and VirtualBox (v7.0) to build a disk image for running experiments in gem5 full-system mode. It also requires Docker Engine (v24+) and Docker Compose (v1.29+) to run all the experiments.

4) *Benchmarks*: Reproducing the performance evaluation results requires SPEC 2017 benchmark [11] and LEBench benchmark [85]. LEBench is included in the artifact source, while SPEC 2017 is not included because of copyright issues.

B. Artifact Installation & Configuration

1) *Pre-setup*: In the artifact folder, run the following command to clone the git submodules:

```
$ git submodule update --init --recursive
```

Note that `cpu2017-1.1.9.iso` is the official image for SPEC 2017 benchmark, which is not released with the artifact since it is not open-source. It should be placed in the artifact folder. The file structure is as shown in Figure 12.




Figure 12: Artifact file structure

2) *Build*: In folder artifact, set up and enter the docker by running:

```
$ cd linux
$ docker-compose up -d
$ docker-compose exec x86_fs /bin/bash
```

In the docker’s shell, build Linux and gem5 by running:

```
$ cd /root/linux
$ python3 compile_scripts/compile.py --num-cores=80
$ cd /root/gem5
$ python3 scripts/compile.py --num-cores=80
```

Please replace 80 with an appropriate number of CPU cores used for the build based on CPU and memory resources.

In the host’s shell and in folder artifact, run the following commands to build the disk image, which takes about 1 hour.

```
$ tar -cvf linux.tar linux
$ cd experiments/disk-image
$ vagrant up
$ vagrant halt
$ vagrant global-status --prune
$ ls ~/VirtualBox\ VMs
$ # Find vm directory denoted as <vm-directory>
$ qemu-img convert -f vmdk -O raw ~/VirtualBox\
  ↳ VMs/<vm-directory>/ubuntu-jammy-22.04-cloudimg.vmdk
  ↳ experiments.img
```

The artifact folder is mounted at /root in the docker. In the following document, we will refer to the path inside the docker by default.

C. Experiment Workflow

For each experiment, we use gem5 to boot Linux under the full system mode and run benchmarks stored in the generated disk image /root/experiments/disk-image/experiments.img. We provide scripts under the /root/gem5/scripts directory to conveniently start experiments using different setups (baseline and *Oreo*) and benchmarks (e.g., SPEC 2017, LEBench, and ASLR attacks).

D. Major Claims

- (C1): *Oreo* introduces negligible overhead compared to the insecure baseline, as demonstrated by experiments E2 and E3 in Figure 8, 9.
- (C2): In SPEC 2017 and LEBench, the ratio of memory accesses where *Oreo* applies its protection (i.e., the ratio of masked memory accesses) is high, which indicates that the benchmarks are relevant to evaluating the performance impacts of *Oreo*. This is demonstrated by experiment E5.
- (C3): *Oreo* prevents the leakage of the ASLR secret offset through paths ① and ②, whereas the insecure baseline does not. This is demonstrated by experiment E4, as shown in Figure 10, 11.

E. Evaluation

1) *Experiment (E1)*: [Functional Test] [1 human-minute + 10 compute-minutes]: This experiment runs two user programs hello and hello_invalid, whose source code can be found in /root/experiments/disk-image/experiments/experiments/src. This experiment demonstrates that *Oreo* allows valid programs (e.g., hello) to run successfully while raising exceptions on invalid memory accesses at commit time (e.g., hello_invalid).

[Preparation] Enter the docker and change the working directory:

```
$ # Host, in folder artifact/linux
$ docker-compose exec x86_fs /bin/bash
$ # Docker
$ cd /root/gem5
```

We will use /root/gem5 as the default working directory for running the following experiments and identifying the path to output files.

[Execution] Running the following command, which takes 10 minutes:

```
$ python3 scripts/gen_checkpoint.py && python3
  ↳ scripts/run_example.py
```

[Results] hello is a valid program, so gem5 should finish simulation without any exceptions. Check result/restore_ko_111_0c0c00/hello_/board.pc.com_1.device for its output. hello_invalid is a malicious program that accesses an invalid address. In the directory result/restore_ko_111_0c0c00/hello_invalid_, check board.pc.com_1.device for its output; check stderr.log for the gem5 exception message on committing invalid memory accesses. The exception message can be located by searching for “ASLR violation” in the log file.

2) *Experiment (E2)*: [Performance Evaluation with SPEC 2017] This experiment evaluates *Oreo*’s performance against the baseline using SPEC 2017 Intrate Benchmarks. In our original setup, we warm up the system and microarchitecture structures by executing 10 billion user-space instructions, and then measure the performance of the next 1 billion user-space instructions. This setup takes several days to finish. For the artifact evaluation, we provide a scaled-down option to warm up with 1 billion user instructions. *Oreo* introduces negligible overhead compared to the baseline.

[Preparation] Same to E1

[Execution] Run the following command:

```
$ python3 scripts/run_spec.py --gen-cpt --begin-cpt=1
  ↳ --num-cpt=1 --num-cores=80 --user-delta=32
  ↳ --spec-inst-warmup-step=1
```

[Results] Run the following command to parse the result:

```
$ python3 scripts/parse_spec.py --parse-raw --begin-cpt=1
  ↳ --num-cpt=1 --roi-idx=2 --expected-stats=3
$ python3 scripts/parse_spec.py --begin-cpt=1 --num-cpt=1
  ↳ --roi-idx=2 --expected-stats=3
```

The CPI overhead introduced by *Oreo* is recorded in scripts/spec_output/merge_input_mean_user_cpi_1_2.pdf. The CPI overhead is expected to be negligible, as shown in Figure 8. However, these benchmarks behave differently in different regions of interest. Hence, the measured absolute CPI might differ from the CPI in Figure 8 when using 1 billion warmup instructions.

3) *Experiment (E3)*: [Performance Evaluation] [1 human-minute + 8 compute-hours]: This experiment evaluates *Oreo*’s performance against the baseline using LEBench, which measures and reports the average latency of various system calls. We ran the benchmark multiple times and used the average latency across all iterations as the final measurement. As shown in Figure 7, *Oreo* introduces negligible average overhead compared to the baseline. The observed variance in latency across different runs is due to the dynamic, non-deterministic scheduling of the Linux kernel, which explains the performance gap between *Oreo* and the baseline in some benchmarks.

[Preparation] Same to E1.

[Execution] We repeated the measurement 16 times to minimize the impact of dynamic scheduling on latency measurement. For artifact evaluation, we offer a scaled-down option to repeat the measurement 8 times by running:

```
$ python3 scripts/run_perf.py --gen-cpt --num-cpt=8
  ↳ --begin-cpt=0 --num-cores=80
```

To modify the number of repetitions, please adjust the `--num-cpt` option. This process takes 8 hours to complete.

[Results] Run the following commands to parse experiment results and generate Figure 7:

```
$ python3 scripts/parse_perf.py --suffix-range=0,8 --plot  
↪ --parse
```

The output files are stored in `/root/gem5/scripts/plot`. Due to Linux’s dynamic scheduling, the overhead measurements may vary from those shown in Figure 7; however, the average overhead is expected to be negligible.

4) *Experiment (E4)*: [Security Evaluation] [5 human-minute + 10 compute-minutes]: This experiment consists of three parts:

- The Prefetch Attack.
- Leakage path ①: we use the speculative code region attack as a representative example of the leakage path ①. We demonstrate that speculatively accessing both valid and invalid addresses produces different microarchitectural traces on the baseline, leading to the leakage of the secret offset. In contrast, *Oreo* produces the same effects regardless of the address validity, confirming that *Oreo* effectively blocks the leakage path ①.
- Leakage path ②: we trigger a system call and record input traces to various microarchitecture structures, which indicates whether the system leaks secrets from the leakage path ②.

[Preparation] Same to E1.

[Execution] Run the following command to run the three parts, which takes about 10 minutes.

```
$ python3 scripts/gen_checkpoint.py && python3  
↪ scripts/run_sec.py
```

[Results]

- The prefetch attack: `scripts/parse_prefetch.py` is used to generate Figure 8 at `scripts/plot/prefetch_plot.pdf`, which indicates that *Oreo* prevents the prefetch attack from leaking secrets through the timing side channel.
- Leakage path ①: For the baseline, the microarchitecture traces of speculatively accessing valid and invalid addresses are at `result/restore_ko_000_0c0c00/blindside_1_0c_/trace.out.gz` and `result/restore_ko_000_0c0c00/blindside_1_0d_/trace.out.gz`. For *Oreo*, the traces are `result/restore_ko_111_0c0c00/blindside_1_0c_/trace.out.gz` and `result/restore_ko_111_0c0c00/blindside_1_0d_/trace.out.gz`. Traces of the baseline are not identical, while traces of *Oreo* are identical, which verifies that *Oreo* blocks the first leakage path. `scripts/parse_trace.py` is used to extract the condensed trace demonstrated in Figure 9 in the directory `scripts/plot/trace` and compare them.
- Leakage path ②: check `iTLBWalker` and `dTLBWalker` traces for both setups (generated by `scripts/parse_trace.py`), which prints the virtual addresses that need to be translated and physical addresses accessed during page table walk. On the baseline, secret dependent physical addresses are used for page table walk, while on

Oreo, secret independent addresses are used. Hence, *Oreo* blocks leakage path ②.

5) *Experiment (E5)*: This experiment reports the ratio of memory accesses where *Oreo* applies its protection (i.e., the masked ratio). This is used to evaluate whether the benchmark is relevant to measure the performance impact of *Oreo*.

[Results] The masked ratio of SPEC 2017 benchmark can be found in the last column of `scripts/spec_output/merge_input_oreo_user_cpi_1_2.csv`. Run the following command, and the masked ratio of LEBench can be found in the last column of `scripts/lebench_output/test_mask_ratio.csv`.

```
$ python3 scripts/parse_perf_stats.py --begin-cpt=0 --num-cpt=8
```

APPENDIX B ATTACK SUMMARY

To complement the attack analysis in Section III, we provide a detailed enumeration of existing microarchitectural-attack-assisted ASLR bypasses, including the leakage paths they take and the utilized side channels in Table I. *Blindside* [27] involves several attacks including Code Region Probing and Spectre Probing. We demonstrated that *Oreo* prevents Code Region Probing in our security evaluation (Section VII-C). Spectre Probing is mitigated by existing Spectre mitigations [5], [6], [8], [15], [20], [23], [39], [40], [44], [48], [53], [58], [63]–[65], [67], [77], [79]–[82]. The Prefetch+Power attack in [49] utilizes power side channels, which are not included in our threat model. *Oreo* successfully mitigates all other attacks in Table I.

Table I: Summary of microarchitectural-attack-assisted ASLR bypasses. The numbers in the second column refer to the three leakage paths in Figure 3.

| Leakage Path | Attacks | Side Channels |
|--------------|--------------------------------------|--|
| ① | Code Region Probing [27] | ICache Prime+Probe |
| | Double Page Fault [33] DrK [37] | Page fault latency to check TLB states |
| | Osiris [76] | Cache side channel to monitor pipeline squashes caused by page fault |
| | EchoLoad [13] | Cache side channel to monitor whether speculation is stalled |
| | Data Bounce [66] Fallout [12] | Cache side channel to monitor store-to-load forwarding behaviors |
| | AMD Prefetch+Time [49] | Prefetch latency to monitor page table walks |
| | AMD Prefetch+Power [49] | Prefetch power consumption to monitor page table walks |
| ② | Gruss et al. [31] EntryBleed [52] | Prefetch latency to check TLB states |
| | TLBleed [28] TagBleed [45] | TLB Prime+Probe |
| | Jump Over ASLR [21] Phantom [78] | BTB Prime+Probe |
| | AnC [29] Binoculars [84] | Cache Prime+Probe to monitor page table walks |
| | Take A Way [50] | DCache Collide+Probe |
| ③ | Spectre Probing [27] | ICache or DCache Prime+Probe |